

Chemical evolution of giant molecular clouds in simulations of galaxies

Alexander J. Richings^{1,2*} and Joop Schaye²

¹*Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA) and Department of Physics and Astronomy, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA*

²*Leiden Observatory, Leiden University, PO Box 9513, 2300 RA Leiden, the Netherlands*

Accepted 2016 May 10. Received 2016 May 09; in original form 2016 April 01

ABSTRACT

We present an analysis of Giant Molecular Clouds (GMCs) within hydrodynamic simulations of isolated, low-mass ($M_* \sim 10^9 M_\odot$) disc galaxies. We study the evolution of molecular abundances and the implications for CO emission and the X_{CO} conversion factor in individual clouds. We define clouds either as regions above a density threshold $n_{\text{H}, \text{min}} = 10 \text{ cm}^{-3}$, or using an observationally motivated CO intensity threshold of 0.25 K km s^{-1} . Our simulations include a non-equilibrium chemical model with 157 species, including 20 molecules. We also investigate the effects of resolution and pressure floors (i.e. Jeans limiters). We find cloud lifetimes up to $\approx 40 \text{ Myr}$, with a median of 13 Myr , in agreement with observations. At one tenth solar metallicity, young clouds ($\lesssim 10 - 15 \text{ Myr}$) are underabundant in H_2 and CO compared to chemical equilibrium, by factors of ≈ 3 and $1 - 2$ orders of magnitude, respectively. At solar metallicity, GMCs reach chemical equilibrium faster (within $\approx 1 \text{ Myr}$). We also compute CO emission from individual clouds. The mean CO intensity, I_{CO} , is strongly suppressed at low dust extinction, A_v , and possibly saturates towards high A_v , in agreement with observations. The $I_{\text{CO}} - A_v$ relation shifts towards higher A_v for higher metallicities and, to a lesser extent, for stronger UV radiation. At one tenth solar metallicity, CO emission is weaker in young clouds ($\lesssim 10 - 15 \text{ Myr}$), consistent with the underabundance of CO. Consequently, X_{CO} decreases by an order of magnitude from 0 to 15 Myr , albeit with a large scatter.

Key words: astrochemistry - molecular processes - ISM: clouds - ISM: molecules - galaxies: ISM.

1 INTRODUCTION

Molecular hydrogen is the main constituent of Giant Molecular Clouds (GMCs), making up most of their mass. However, cold H_2 is difficult to observe in emission, as the lowest rotational transition of the H_2 molecule has an excitation energy of $E/k_B = 510 \text{ K}$ (Dabrowski 1984). It is therefore difficult to excite H_2 at the cold temperatures typical of GMCs ($\sim 10 \text{ K}$).

CO is typically the next most abundant molecule in GMCs. It is also much easier to excite the rotational and vibrational levels of the CO molecule at low temperatures. For example, the lowest rotational transition of CO ($J = 1 - 0$) has an excitation energy of $E/k_B = 5.53 \text{ K}$. CO emission is therefore commonly used as a tracer of molecular gas in GMCs (e.g. Solomon et al. 1987; Dame et al. 2001; Heyer et

al. 2001). The velocity-integrated CO intensity, I_{CO} , is then converted to an H_2 column density, N_{H_2} , using a conversion factor X_{CO} , defined as:

$$X_{\text{CO}} = \frac{N_{\text{H}_2}}{I_{\text{CO}}} \text{ cm}^{-2} (\text{K km s}^{-1})^{-1}. \quad (1.1)$$

To accurately determine the molecular content of a GMC in this way, we therefore require a detailed understanding of the X_{CO} factor, including how it depends on the physical conditions in the GMC, such as its metallicity and the radiation field.

There have been many studies, both observational and theoretical, to determine the X_{CO} factor (see Bolatto et al. 2013 for a recent review). Observational studies use various methods to determine the total molecular content, which can then be compared to the CO emission to determine the X_{CO} factor. For example, virial techniques assume that the GMC

* Email: a.j.richings@northwestern.edu

is in virial equilibrium, which allows one to measure the total mass of a GMC from its size and velocity dispersion, which is assumed to be the molecular mass (e.g. [Scoville et al. 1987](#); [Solomon et al. 1987](#)). Other studies estimate the dust content of GMCs, either by mapping the extinction towards background stars (e.g. [Frerking et al. 1982](#); [Lombardi et al. 2006](#); [Pineda et al. 2008](#)), or by measuring dust emission in the far-infrared (e.g. [Dame et al. 2001](#); [Planck Collaboration XIX 2011](#)). This can then be converted into a total gas column density, assuming a dust-to-gas ratio. Diffuse gamma-ray emission arising from interactions between cosmic rays and nucleons can also be used to estimate the total gas column density (e.g. [Strong & Mattox 1996](#); [Abdo et al. 2010](#); [Ackermann et al. 2012](#)).

Some theoretical studies of the X_{CO} factor use models of photodissociation regions (PDRs), where a cloud of gas is illuminated from one side by an external UV radiation field. [Tielens & Hollenbach \(1985\)](#) use PDR models to determine the chemical and temperature structure of such clouds for various gas densities and radiation fields. [van Dishoeck & Black \(1988\)](#) and [Visser et al. \(2009\)](#) focus on the chemistry and photodissociation of CO in PDR models, and they use these models to determine how the CO column density varies with dust extinction. [Sternberg et al. \(2014\)](#) recently presented a detailed study of the H I-to-H₂ transition in clouds, using both analytic theory and numerical PDR models. These PDR models assume that the abundances of molecules and atoms are in chemical equilibrium, or a ‘steady state’, and that the clouds have a constant density profile.

These PDR models can then be used to study how the X_{CO} factor depends on the physical conditions. For example, [Bell et al. \(2006\)](#) use PDR models to explore how X_{CO} varies in different environments. They find that, at low dust extinction, A_V , X_{CO} decreases with increasing A_V , until it reaches a minimum and subsequently increases with A_V once the CO line becomes optically thick. They show that the $X_{\text{CO}} - A_V$ profile depends on cloud properties, including gas density, radiation field strength, metallicity and turbulent velocity dispersion.

Other theoretical studies of the X_{CO} factor use hydrodynamic simulations of a turbulent interstellar medium (ISM) to study the environmental dependence of the X_{CO} factor, which account for more realistic cloud geometries (e.g. [Glover & Mac Low 2011](#); [Shetty et al. 2011a,b](#); [Clark & Glover 2015](#)). [Narayanan et al. \(2011, 2012\)](#) combine hydrodynamic simulations of isolated and merging galaxies, using a subgrid model for cold gas below 10^4 K, with radiative transfer calculations of dust and molecular line emission to explore how galaxy mergers and the galactic environment affect the X_{CO} factor. [Feldmann et al. \(2012\)](#) combine the results of sub-parsec resolution simulations from [Glover & Mac Low \(2011\)](#) with gas distributions from the cosmological simulations of [Gnedin & Kravtsov \(2011\)](#) to model the X_{CO} factor, finding a metallicity dependence of X_{CO} (averaged over kpc scales) of $X_{\text{CO}} \propto Z^{-\gamma}$, where $\gamma \approx 0.5 - 0.8$.

Theoretical models of X_{CO} need to determine the abundances of CO and H₂ under various conditions. The simplest approach is to assume that these abundances are in chemical equilibrium (e.g. [Narayanan et al. 2011, 2012](#)). However, this assumption may not be valid if the formation time-scale of molecules is comparable to the lifetimes of GMCs, par-

ticularly in young clouds. Observational estimates have suggested a wide range of GMC lifetimes, from a few Myr (e.g. [Elmegreen 2000](#)), to $\approx 20 - 40$ Myr (e.g. [Bash et al. 1977](#); [Kawamura et al. 2009](#); [Murray 2011](#); [Miura et al. 2012](#)), to hundreds of Myr (e.g. [Scoville et al. 1979](#)).

[Bell et al. \(2006\)](#) include time-dependent chemistry of H₂ and CO in their PDR models, with metallicities $0.01 \leq Z/Z_{\odot} \leq 1.0$ ¹, and they consider various cloud ages. They find significant evolution in the X_{CO} factor at times $\lesssim 1$ Myr, with less evolution for cloud ages $1 - 10$ Myr, and no notable evolution beyond 10 Myr, even though it takes up to 100 Myr for the chemical abundances to reach steady-state in their models. [Glover & Mac Low \(2011\)](#) and [Shetty et al. \(2011a,b\)](#) also include time-dependent chemistry in their simulations of a turbulent ISM, with metallicities $0.03 \leq Z/Z_{\odot} \leq 1.0$ and $0.1 \leq Z/Z_{\odot} \leq 1.0$, respectively. However, since they include only a region of the ISM in their simulations, and not an entire galaxy, they may be missing some aspects of the evolution of GMCs in a galactic environment. Indeed, [Dobbs & Pringle \(2013\)](#) explore GMC evolution in simulations of isolated disc galaxies, with solar metallicity, and they find complex evolutionary histories. GMCs in their simulations often form by assembling from smaller clouds and ambient ISM material, or by breaking off from larger clouds, while they are dispersed by stellar feedback and shear, or are accreted onto larger clouds. It would therefore be useful to explore the chemical evolution of GMCs within a realistic galactic environment.

In this paper we investigate how the molecular abundances of GMCs evolve, and under what conditions these abundances are out of chemical equilibrium. We consider the effects of cloud age, metallicity and the radiation field. We can then determine how the conditions affect the X_{CO} factor. We study clouds of dense gas ($n_{\text{H}} > 10 \text{ cm}^{-3}$) in the high-resolution Smoothed Particle Hydrodynamics (SPH) simulations of isolated disc galaxies presented in [Richings & Schaye \(2016\)](#), hereafter Paper I. These simulations include a treatment for the non-equilibrium chemistry of 157 species, including 20 different molecules ([Richings et al. 2014a,b](#)). We also run radiative transfer calculations on these simulations in post-processing to determine the $^{12}\text{CO } J = 1 - 0$ line emission² from individual GMCs, and hence compute their X_{CO} factors.

The remainder of this paper is organised as follows. In section 2 we summarise the simulations and initial conditions from paper I. In section 3 we describe the methods that we use to analyse GMCs in these simulations, including how we identify clouds, how we link clouds in previous and subsequent snapshots to identify their progenitors and descendants, and how we create maps of CO emission from individual clouds in post-processing. In section 4 we investigate the scaling relations of these clouds and compare them to observations. In section 5 we look at the H₂ and CO abundances of our simulated GMCs as a function of cloud age to explore their chemical evolution. In section 6 we use the CO

¹ Throughout this paper we use a solar metallicity of $Z_{\odot} = 0.0129$ ([Wiersma et al. 2009](#)), although other studies that we quote in this paragraph use different definitions of Z_{\odot} .

² For the remainder of this paper, we will use ‘CO’ to refer to ^{12}CO , unless stated otherwise.

$J = 1 - 0$ line emission from simulated clouds to investigate the X_{CO} factor, and we summarise our main results in section 7. Finally, in Appendix A we explore how our results are affected by changing the resolution of our simulations, and in Appendix B we explore how our results are affected by the pressure floor that we impose in our simulations to ensure that the Jeans mass is always well-resolved.

2 SIMULATIONS

We study GMCs in the suite of hydrodynamic simulations of isolated disc galaxies that were first presented in paper I. The details of how these simulations were run, along with properties of the galaxies such as their star formation histories and their outflow rates and velocities, can be found in paper I. Here we summarise the main features of these simulations.

The simulations were run using a modified version of the tree/SPH code GADGET3, last described in Springel (2005). The hydrodynamics solver has been replaced with the suite of hydrodynamical methods collectively known as ANARCHY, which incorporates many of the latest improvements on ‘classical’ SPH methods, including the pressure-entropy formulation of SPH, as derived by Hopkins (2013); a switch for artificial conduction, similar to the one used by Price (2008); a switch for artificial viscosity, from Cullen & Dehnen (2010); the time-step limiters from Durier & Dalla Vecchia (2012); and the C^2 Wendland (1995) kernel, for which we use 100 neighbours. ANARCHY will be described in more detail in Dalla Vecchia (in preparation); see also Appendix A of Schaye et al. (2015) for a full description of our version of ANARCHY.

2.1 Chemistry and subgrid models

We follow the chemical evolution of the abundances of ions and molecules in the gas using the chemical model of Richings et al. (2014a,b). This model includes all ionisation states of the 11 elements that contribute most to the cooling rate³, along with 20 molecular species⁴, most importantly H_2 and CO. This gives us a chemical network of 157 species in total. The chemical species evolve via collisional ionisation, radiative and di-electronic recombination, charge transfer reactions, photoionisation (including Auger ionisation), cosmic ray ionisation (parameterised by an HI cosmic ray ionisation rate of $2.5 \times 10^{-17} \text{ s}^{-1}$; Williams et al. 1998), and various molecular reactions, including the formation of H_2 on dust grains (Cazaux & Tielens 2002) and in the gas phase.

The photoionisation, photoheating and photoelectric dust heating rates are computed assuming a constant, uniform UV radiation field, either the local interstellar radiation field (ISRF) of Black (1987), or ten per cent of this ISRF. We also use a self-shielding prescription to account for the attenuation of photochemical rates by dust and gas (Richings et al. 2014b). This prescription includes self-shielding of H_2 and CO, and shielding of CO by H_2 . We assume that

the shielding occurs locally, which allows us to express the column density of each particle as the density, ρ , multiplied by a local shielding length, L . For the shielding length, we use a local Sobolev-like approximation, $L = \rho/|2\nabla\rho|$ (e.g. Gnedin et al. 2009).

From the chemical network we obtain a system of 158 differential equations (157 chemical rate equations and the thermal equation for the temperature evolution), which we integrate for each gas particle over each hydrodynamic time-step using the implicit differential equation solver CVODE, from the SUNDIALS⁵ suite. This enables us to follow the non-equilibrium evolution of ion and molecule abundances, and also to evolve the temperature using cooling rates computed from these abundances, without needing to assume chemical equilibrium.

Gas particles are allowed to form stars if their hydrogen number density, n_{H} , exceeds a threshold of 1 cm^{-3} and their temperature is below 1000 K. If a particle meets these criteria, it forms stars at a rate per unit volume given by the gas density over the local free fall time, multiplied by an efficiency factor ϵ_{SF} , which we take to be 0.005. Gas particles are then stochastically converted into star particles according to a probability that is determined from the particle’s star formation rate and the hydrodynamic time-step. The value of the efficiency, ϵ_{SF} (and the value of the heating temperature, ΔT , used in the stellar feedback model; see below) were chosen to reproduce the observed Kennicutt-Schmidt relation (see fig. 3 in paper I).

We include feedback from star formation using a thermal supernova prescription similar to that of Dalla Vecchia & Schaye (2012), with some modifications. As each star particle is treated as a simple stellar population (rather than an individual star), we can calculate the number of supernovae that explode from each star particle in a given time-step, using the stellar lifetimes of Portinari et al. (1998) and assuming a Chabrier (2003) initial mass function (IMF). When supernovae explode, their energy is injected into the gas thermally by stochastically selecting neighbouring gas particles to be heated by $\Delta T = 10^{7.5} \text{ K}$.

By imposing a minimum heating temperature, we ensure that we reduce artificial radiative losses due to our finite resolution, which might otherwise make the stellar feedback unrealistically inefficient. We are unable to resolve individual supernovae. Instead, at the fiducial resolution used in our simulations (750 M_{\odot} per gas particle) each heating event corresponds to approximately ten supernovae exploding simultaneously. The probability of stochastically selecting a gas particle to be heated is computed such that, when averaged over time and over all particles in the simulation, the expectation value for the total injected thermal energy is equal to the total available energy from supernovae.

The difference between our stellar feedback model and that of Dalla Vecchia & Schaye (2012) is that we distribute the total available supernova energy from each star particle over time, according to the lifetimes of massive stars, rather than injecting it all at 30 Myr after the birth of the star particle.

To ensure that the Jeans mass is always resolved, we impose a density-dependent pressure floor, P_{floor} , in the hy-

³ H, He, C, N, O, Ne, Mg, Si, S, Ca, Fe

⁴ H_2 , H_2^+ , H_3^+ , OH, H_2O , C_2 , O_2 , HCO^+ , CH, CH_2 , CH_3^+ , CO, CH^+ , CH_2^+ , OH^+ , H_2O^+ , H_3O^+ , CO^+ , HOC^+ , O_2^+

⁵ <https://computation.llnl.gov/casc/sundials/main.html>

Table 1. Properties of the galaxy simulations used in this paper: total mass M_{200} within the radius $R_{200,\text{crit}}$ enclosing a mean density of 200 times the critical density of the Universe at redshift zero, NFW concentration c_{200} of the dark matter halo, initial stellar mass $M_{*,\text{init}}$, initial gas mass $M_{\text{gas,init}}$, disc gas mass fraction $f_{\text{d,gas}}$, mass per gas or star particle m_{baryon} , gravitational softening length ϵ_{soft} , gas metallicity Z , and UV radiation field. Parameters highlighted in bold have been changed from the ref model.

Model	M_{200} (M_{\odot})	c_{200}	$M_{*,\text{init}}$ (M_{\odot})	$M_{\text{gas,init}}$ (M_{\odot})	$f_{\text{d,gas}}$	m_{baryon} (M_{\odot})	ϵ_{soft} (pc)	Z (Z_{\odot})	UV Field
ref	10^{11}	8.0	1.4×10^9	4.8×10^8	0.3	750	3.1	0.1	ISRF ^a
hiZ	10^{11}	8.0	1.4×10^9	4.8×10^8	0.3	750	3.1	1.0	ISRF
lowISRF	10^{11}	8.0	1.4×10^9	4.8×10^8	0.3	750	3.1	0.1	10% ISRF

^a ISRF of [Black \(1987\)](#)

hydrodynamic equations, such that the Jeans mass will always be at least a factor $N_{\text{J,m}}$ times the mass within the SPH kernel. This is similar to the methods used by e.g. [Robertson & Kravtsov \(2008\)](#); [Schaye & Dalla Vecchia \(2008\)](#); [Hopkins et al. \(2011\)](#). The pressure floor is given by equation 2.12 of paper I:

$$P_{\text{floor,m}} = \left(\frac{36}{\pi^5} \right)^{1/3} \frac{G}{\gamma} (N_{\text{J,m}} N_{\text{ngb}}^{\text{SPH}} m_{\text{gas}})^{2/3} \rho^{4/3}, \quad (2.1)$$

where $\gamma = 5/3$ is the ratio of specific heats, $N_{\text{ngb}}^{\text{SPH}}$ is the number of SPH neighbours, m_{gas} is the mass per SPH particle, and ρ is the gas density. We use a conservative fiducial value of $N_{\text{J,m}} = 4$ in our simulations, but see Appendix B for the effects of lowering this pressure floor. We impose this Jeans limiter as a pressure floor rather than a temperature floor (as used by [Schaye & Dalla Vecchia 2008](#)) so that gas particles can continue to cool below the temperature corresponding to the pressure floor, and thus will evolve towards thermal and chemical equilibrium for the given density.

2.2 Initial conditions

We ran simulations of isolated disc galaxies using initial conditions based on the model of [Springel et al. \(2005\)](#). These initial conditions were generated using a modified version of a code that was kindly provided to us by Volker Springel. Each galaxy has a total mass within $R_{200,\text{crit}}$ (i.e. the radius enclosing 200 times the critical density) of $M_{200} = 10^{11} M_{\odot}$. The galaxies initially consist of a rotating disc of gas and stars and a central stellar bulge, embedded in a dark matter halo. The initial stellar mass is $M_* = 1.4 \times 10^9 M_{\odot}$, which is consistent with the abundance matching results of [Moster et al. \(2013\)](#) corrected for baryonic effects according to the prescription of [Sawala et al. \(2015\)](#). Twenty per cent of the initial stellar mass is in the bulge, with the remainder in the stellar disc. We use a gas mass fraction in the disc of 30 per cent, which gives an initial gas mass of $4.8 \times 10^8 M_{\odot}$.

The gas and stellar discs initially have an exponential surface density profile with a radial scale length of 2.0 kpc. The vertical structure of the stellar disc has an isothermal profile with a scale height of ten per cent of the radial scale length, while the gas is initially in chemical equilibrium with a constant temperature of 10^4 K and a vertical structure set up in hydrostatic equilibrium using an iterative procedure. At this temperature, most of the hydrogen is in HII

in chemical equilibrium. The stellar bulge has a [Hernquist \(1990\)](#) density profile, and the dark matter halo follows a [Hernquist \(1990\)](#) profile that is scaled to match a [Navarro et al. \(1996\)](#) (NFW) profile in the inner regions with a concentration $c_{200} = 8.0$, which agrees with the redshift zero mass-concentration relation of [Duffy et al. \(2008\)](#).

We use a fiducial resolution of $750 M_{\odot}$ per gas or star particle, with 100 SPH neighbours, and a gravitational softening length of 3.1 pc (but see Appendix A for runs with a factor four higher/lower mass resolution), and we model the dark matter halo using a static potential. Each simulation initially contains 6.45×10^5 gas particles and 1.88×10^6 star particles.

We include a constant, uniform UV radiation field, along with a local self-shielding prescription, and the gas metallicity is held fixed, with dust-to-gas mass ratios of $2.4 \times 10^{-3} Z/Z_{\odot}$ and $4.0 \times 10^{-3} Z/Z_{\odot}$ for graphite and silicate grain species, respectively. These dust-to-gas ratios were taken from the ‘ISM’ grain abundances used by the photoionisation code CLOUDY⁶ version 13.01 ([Ferland et al. 2013](#)), and we assume that they scale linearly with metallicity, Z . However, this assumption of a linear scaling between dust-to-gas ratio and metallicity may not be accurate, particularly at low metallicity. For example, the dust content of low-metallicity dwarf galaxies is found to be less than what one would expect from a linear scaling (e.g. [Rémy-Ruyer et al. 2014](#)).

In paper I, we ran six simulations with different combinations of metallicity and UV radiation field. Each simulation was repeated twice, once with the full non-equilibrium chemical model of [Richings et al. \(2014a,b\)](#), and once using cooling rates computed assuming chemical equilibrium. In this paper, we focus on three of these simulations: ref (ten per cent solar metallicity and the local ISRF of [Black 1987](#)), hiZ (solar metallicity and the [Black 1987](#) ISRF), and low-ISRF (ten per cent solar metallicity and ten per cent of the [Black 1987](#) ISRF), all evolved using the full non-equilibrium chemical model. We focus on these as they are the most relevant for conditions in molecular clouds in low-mass galaxies. Of the three remaining simulations in paper I, lowZ (with one per cent solar metallicity) did not form dense clouds, as the gas was mostly unable to cool to a cold (~ 100 K) phase; UVB (evolved with the redshift zero UV background

⁶ <http://nublado.org/>

of Haardt & Madau 2001) used an extragalactic UV radiation field that is more relevant for the circum- and intergalactic medium than for molecular clouds; and UVBthin neglected self-shielding of UV radiation, which is necessary for the formation of molecules. The properties of our simulations are summarised in Table 1.

Our simulations all use a constant HI cosmic ray ionisation rate of $2.5 \times 10^{-17} \text{ s}^{-1}$ (Williams et al. 1998). By keeping the cosmic ray rate fixed as we vary the strength of the UV radiation field, we can isolate the effects of the UV radiation alone on the molecular clouds. However, in reality, it is likely that both the UV radiation and the cosmic rays are produced in star forming regions. Therefore, we would expect both to vary in proportion to the local star formation rate (see e.g. Clark & Glover 2015 for examples of varying both the UV radiation field and the cosmic ray ionisation rate together in simulations of molecular clouds).

3 ANALYSIS METHODS

In this section we describe the methods that we use to analyse gas clouds in our simulations, including the algorithm that we use to identify clouds (§3.1), how we link clouds to their progenitors and descendants to define their mass evolution (§3.2), and how we create maps of CO emission from individual clouds (§3.3).

3.1 Clump finding algorithm

Observationally, molecular clouds are typically identified as regions detected in emission from a molecular tracer (often CO) above an intensity threshold (e.g. Larson 1981; Solomon et al. 1987). This is approximately equivalent to selecting regions above a molecular gas surface density threshold. However, as we are interested in the atomic to molecular transition, we do not want to select only clouds that are already molecular. We therefore need a criterion that is based on the total gas content, and not just on the molecular content.

Furthermore, Dobbs et al. (2015) found that using a grid-based approach to identify clouds above a surface density threshold in simulations can create problems for studying the cloud evolution. They found that clouds that are identified with such a method appear to evolve on shorter time-scales than is seen in the three-dimensional particle distribution. These errors arise due to the projection onto a two-dimensional grid, as the gas moves relative to the grid.

We therefore base our clump finding algorithm on the particle-based approach used by Dobbs et al. (2015). This is a Friends-of-Friends (FoF) algorithm that acts on dense gas particles. We first select gas particles with a hydrogen number density, n_{H} , above a threshold $n_{\text{H}, \text{min}}$. We then link together nearby dense particles by taking each particle in turn and identifying particles that lie within a linking length, l .

There are two parameters in this method, $n_{\text{H}, \text{min}}$ and l . However, as noted by Dobbs et al. (2015), they are degenerate, as denser particles will be closer together. We use a density threshold of $n_{\text{H}, \text{min}} = 10 \text{ cm}^{-3}$, which is comparable to the density at which we expect the transition from atomic to molecular hydrogen to occur (e.g. Schaye 2001; Gnedin et al. 2009). This ensures that we focus on clouds that are likely




Figure 1. Maps of gas surface density after 500 Myr from simulations lowISRF ($0.1 Z_{\odot}$, ten per cent of the Black (1987) ISRF; top), ref ($0.1 Z_{\odot}$, Black (1987) ISRF; centre) and hiZ (Z_{\odot} , Black (1987) ISRF; bottom). Each map is 8 kpc across and views the disc of the galaxy face-on. We also zoom in (by a factor 13) on the most massive clouds in each simulation. We see a wide range of morphologies, including nearly spherical clouds, clouds that have been sheared into long filaments, and clouds with multiple density peaks that are indicative of cloud mergers.




Figure 2. Cloud mass functions from simulations lowISRF (black solid curve), ref (red dashed curve) and hiZ (blue dot-dashed curve), taken from snapshots at 100 Myr intervals from 100 Myr to 900 Myr. The vertical dotted line shows a mass of $3.75 \times 10^4 M_\odot$, which corresponds to the mass of a cloud containing 50 particles. We consider only clouds above this mass for the remainder of this study. All three simulations show similar cloud mass functions.

to become molecular. We then use a linking length $l = 10$ pc, which corresponds approximately to the mean spacing between gas particles at the density threshold, $n_{\text{H}, \text{min}}$, for our resolution of $750 M_\odot$ per particle.

Fig. 1 shows maps of the gas surface density in each of our three simulations after 500 Myr. Each map is 8 kpc across and views the disc face-on. We also zoom in (by a factor 13) on the six most massive clouds in each simulation. We see that these clouds show a wide range of morphologies. Some are approximately spherical, while others have been stretched into long, thin filaments by shear in the rotating disc. We also see some clouds with two or more density peaks, which suggests that they consist of multiple clumps that are in the process of merging.

Fig. 2 shows the cloud mass functions for each simulation. Here we have identified clouds in snapshots at 100 Myr intervals, from 100 Myr to 900 Myr, and combined the different snapshots into a single mass function for each simulation. We will show in the next section that the clouds have lifetimes < 100 Myr if we define lifetimes using the particles originally in the cloud when it is identified, so we do not double-count cloud mass by combining snapshots in this way. However, if we follow the total mass of cloud progenitors/descendants, we find that individual cloud structures can survive for > 100 Myr, although new gas has cycled through them. Such long-lived cloud structures will appear multiple times in Fig. 2.

All three simulations show similar cloud mass functions. For the remainder of this study, we shall focus on clouds that contain at least 50 gas particles to avoid poorly resolved clouds in our analysis (but see Appendix B for the effects of the pressure floor on low-mass clouds). This corresponds to

a mass of $3.75 \times 10^4 M_\odot$, shown by the vertical dotted line in Fig. 2.

We also need to define the radius of each cloud, which will be important for comparing to the observed molecular cloud scaling relations (see §4). We determine the radius by finding the 3-dimensional ellipsoid that approximately encloses the particles in the cloud. First, we compute the moment of inertia tensor, I :

$$I_{ij} = \sum_{k=1}^N m_k (|\mathbf{r}|^2 \delta_{ij} - r_{k,i} r_{k,j}), \quad (3.1)$$

where m_k is the mass of the k^{th} particle, \mathbf{r}_k is the position vector of the k^{th} particle in the cloud's centre of mass frame, the summation is over the N particles in the cloud, i and j index the Cartesian directions ($i, j = 1, 2, 3$ in 3d), and δ_{ij} is the Kronecker delta function. The eigenvectors of I give the directions of the principle axes of the cloud. We then determine the maximum extent of the particle distribution along each principle axis to obtain the semi-major, intermediate and minor axes, a , b and c respectively, of the ellipsoid that approximately encloses the particles in the cloud. Finally, we define the cloud radius, R_{mean} , to be the geometric mean of these three axes, i.e.:

$$R_{\text{mean}} = (abc)^{1/3}. \quad (3.2)$$

The above cloud definition is based on a density threshold. However, observations define molecular clouds based on a CO intensity threshold. We therefore also consider an alternative cloud definition, based on the CO emission, which we discuss in section 3.3.

3.2 Cloud mass evolution

To follow the mass evolution of individual clouds, and hence determine their ages and lifetimes, we first ran the clump finding algorithm described above on all snapshots, taken at intervals of 1 Myr. We then took each massive cloud (containing at least 50 particles) in a given snapshot and traced back its main progenitor in preceding snapshots, and its main descendant in subsequent snapshots.

There are a couple of different ways in which we can link a given cloud to its progenitors and descendants. In the first method that we use, we first take all particles in a given cloud in snapshot i . We then look for these particles in the preceding snapshot, $i - 1$, and we identify the cloud that contains the most of these particles. This cloud is selected as the main progenitor. We then take all of the particles in the main progenitor, and we look for these particles in snapshot $i - 2$. We repeat this process to find the main progenitor in each preceding snapshot until we can no longer identify a progenitor. Finally, we repeat the above procedure in the snapshots following i to identify the main descendants.

The above method allows us to follow the evolution of the total mass of the cloud. In this way, we can trace coherent cloud structures through time. However, gas will cycle through individual clouds, with new gas being added to the cloud via smooth accretion or mergers, while existing gas can break off into smaller clouds or disperse into the ISM. Therefore, after some time, it is possible that a cloud will




Figure 3. Mass evolution of four clouds selected at 500 Myr from the ref simulation. We show the evolution of the total cloud mass (black solid curves) and of the mass of particles originally in the cloud at 500 Myr (blue dashed curves). The horizontal dotted lines indicate half of the mass of the cloud at 500 Myr, which we use to define the cloud ages and lifetimes (see text).

no longer contain any of the material that was originally in the cloud in snapshot i .

We therefore also considered an alternative method to link clouds with their progenitors and descendants, in which we consider only gas particles that were originally in the cloud in snapshot i . This is similar to how [Dobbs & Pringle \(2013\)](#) trace the evolution of GMCs in their simulations of an isolated disc galaxy. We first take the particles in the cloud in snapshot i and identify the main progenitor in snapshot $i-1$ that contains the most of these particles, as before. However, we then take only the particles in the main progenitor that were originally in the cloud in snapshot i (and not all of the main progenitor’s particles), and we trace these back to snapshot $i-2$ to find the preceding main progenitor, and so on. We then repeat this procedure in later snapshots to identify the main descendants. In this way we can trace the evolution of only gas that was originally in the cloud in snapshot i .

In Fig. 3 we show examples of the mass evolution of individual clouds selected at 500 Myr from the ref simulation. For each cloud we show the evolution of the total cloud mass (black solid curves), using the first method described above, and of the mass of original particles that were in the cloud at 500 Myr (blue dashed curves), using the second method described above. The horizontal dotted line in each panel indicates half of the mass of the cloud at 500 Myr, which we use to define the age and lifetime of the cloud (see below).

For some clouds, the evolution of the total and original mass are similar (e.g. the top right panel). These are clouds that reach their peak mass close to 500 Myr, and have fairly simple evolutionary histories, for example with no significant cloud mergers bringing in new material at later times.

In many other clouds, the evolution is more complex. For example, in the bottom right panel, the cloud is still




Figure 4. Cloud ages (black circles), defined as the time since the cloud’s main progenitor had half of its current mass, and cloud lifetimes (red squares), defined as the period from when the main progenitor had half of its current mass to when the main descendant is reduced to half of its current mass. We define ages and lifetimes using either the total progenitor/descendant mass (top panel), or the mass of particles that were originally part of the cloud, i.e. at the time the cloud was identified (bottom panel). We show all clouds with at least 50 particles, selected in snapshots at 100 Myr intervals from 100 Myr to 900 Myr, from the ref simulation. Our other two simulations (lowISRF and hiZ; not shown) have similar distributions of cloud ages and lifetimes. The horizontal stripes of points at low ages/lifetimes are due to the finite interval (1 Myr) between snapshots.

growing at 500 Myr. The total mass of the cloud therefore quadruples over the following 19 Myr, after which it rapidly declines. However, by definition, the original mass is a maximum in the original snapshot (at 500 Myr in this example). We see that the original mass in this example remains nearly constant over the same period.

Finally, in some clouds (e.g. the top left panel), we find that the progenitors and descendants traced by the total mass extend over a much longer time period than those traced only by the original particles at 500 Myr. These are clouds that are constantly cycling through new gas, via accretion and cloud mergers, while existing gas breaks away or is blown away and disperses. The cloud in the top left panel is located at the centre of the galaxy. We saw in Fig. 1 that there is more dense gas near the centre, with several clouds packed closely together within the central few hundred parsecs. This explains why we see a strong cycling of gas through individual clouds in this region.

We can now use the mass evolution of a cloud’s progenitors and descendants to determine the age and lifetime of the cloud, based on either the total mass of the cloud or the

mass of original particles. We define the age of the cloud as the time since the mass was half of its current value, and we define its lifetime to be the total period over which its mass is greater than half of its current value. For example, suppose we identify a cloud at time t_{now} . In the past, its main progenitor had half of its current mass at time t_{past} , and in the future, its main descendant is reduced to half of its current mass at time t_{future} . The age is then $t_{\text{now}} - t_{\text{past}}$, and the lifetime is $t_{\text{future}} - t_{\text{past}}$. The ages and lifetimes will depend on the mass fraction that we use to define them. For example, if we use the time when the main progenitor/descendant was a quarter of the cloud's current mass, rather than half, the median lifetimes are increased by ≈ 50 per cent. However, by using a factor of half, there can only be one 'main' progenitor/descendant in each snapshot over the cloud's lifetime, and we avoid ambiguities arising from multiple progenitors/descendants with equal mass.

In Fig. 4 we plot the ages (black circles) and lifetimes (red squares) of clouds from the ref simulation versus their current mass, using the total progenitor/descendant mass (top panel) or the mass of original particles (bottom panel). Our other two simulations (lowISRF and hiZ; not shown) have similar distributions of cloud ages and lifetimes. We show all clouds with at least 50 particles identified in snapshots at 100 Myr intervals, from 100 Myr to 900 Myr. Note that, while we only show clouds with at least 50 particles, we still identify clouds with as few as 25 particles, so we can trace the clouds in Fig. 4 to the time when they had, or will have, half of their current mass.

If we use the total cloud mass (top panel), we find a median cloud age of 12 Myr, and a median lifetime of 33 Myr. There is a lot of scatter in cloud ages and lifetimes, with many having ages and lifetimes of a few hundred Myr. Note that, since we combine snapshots at 100 Myr intervals in this figure, evolutionary tracks will appear multiple times if they have a lifetime longer than this interval.

In the bottom panel, we see that cloud ages and lifetimes defined using only the original particles are shorter than those defined from the total mass. Using this definition, we find a median age of 5 Myr and a median lifetime of 13 Myr. There is again a lot of scatter in cloud ages and lifetimes, but we find that most clouds have an age $\lesssim 30$ Myr and a lifetime $\lesssim 40$ Myr. Observational estimates, typically based on associated signatures of star formation such as young stellar clusters and HII regions, find GMC lifetimes $\approx 20 - 40$ Myr (e.g. Bash et al. 1977; Kawamura et al. 2009; Murray 2011; Miura et al. 2012), although Elmegreen (2000) and Scoville et al. (1979) find lifetimes of a few Myr and hundreds of Myr, respectively. We find no clear trend of age or lifetime with the current mass of the cloud.

Since we run each simulation for 1 Gyr, we follow the evolution of the galaxy for many cloud lifetimes. This is important as it ensures that the evolution of individual clouds is not strongly affected by the initial chemical state of the gas at the beginning of the simulation, when most of the hydrogen was in HII.

An important caveat to note is that the stellar feedback model used in these simulations is for feedback from supernovae. This means that we do not explicitly model feedback processes that act on shorter time-scales than supernovae, such as stellar winds and photoheating from HII regions.

Such processes might disrupt clouds before the supernovae explode, thus shortening their ages and lifetimes.

For the remainder of this paper, we will use cloud ages and lifetimes defined via the mass of original particles (i.e. the bottom panel of Fig. 4). This definition gives a better indication of how long the current material has been in the cloud. However, both age/lifetime definitions that we have considered here (using total or original mass) involve tracing individual particles through time in the simulations, which is not possible in observations. Observational estimates are typically based on nearby signatures of star formation, such as young stellar clusters and HII regions (e.g. Kawamura et al. 2009). It is not clear which of our definitions is likely to correspond more closely with these observational definitions, so we need to be careful when comparing to observed GMC lifetimes.

3.3 CO emission maps

We computed CO emission from the $J = 1 - 0$ line in our simulations in post-processing, using the publicly available Monte-Carlo radiative transfer code RADMC-3D⁷ (version 0.38), written by Cornelis Dullemond. This code follows emission from user-specified molecular and atomic lines, and also includes thermal emission, absorption and scattering from dust grains. We used molecular CO data from the LAMDA database⁸ (Schöier et al. 2005), including collisional excitation rates of CO by ortho- and para-H₂ (Yang et al. 2010). We assumed an ortho-to-para ratio of 3:1 for H₂. We included two species of dust grains, graphite and silicate, with dust opacities from Martin & Whittet (1990), who used the power-law size distribution of dust grains from Mathis et al. (1977).

Line emission from CO depends on the level populations of the CO molecule. The simplest method is to assume that the level populations are in Local Thermodynamic Equilibrium (LTE). However, this assumption may not always be valid. We therefore computed the level populations in non-LTE using the Local Velocity Gradient (LVG) method, also known as the Sobolev (1957) approximation. This method assumes that, due to gas motions, photons emitted from transitions in the CO molecule will become sufficiently Doppler shifted after travelling some distance that the photon can no longer be absorbed by the same transition that produced it. This allows us to define an escape probability for these photons based on their velocity gradient. We can then determine the level populations, including radiative excitation by line photons, from local quantities alone. A more detailed description of the LVG method, as implemented in RADMC-3D, can be found in Shetty et al. (2011a).

In addition to thermal broadening of the emission line, RADMC-3D also allows the inclusion of doppler broadening by unresolved microturbulence. In our simulations, we impose a density-dependent pressure floor, P_{floor} , on the gas to ensure that the Jeans mass is resolved by at least 4 SPH kernel masses, to prevent artificial fragmentation

⁷ <http://www.ita.uni-heidelberg.de/~dullemond/software/radmc-3d/>

⁸ <http://home.strw.leidenuniv.nl/~moldata/>




Figure 5. CO $J = 1-0$ line emission maps (*top row*) and gas surface density maps (*bottom row*) of molecular clouds from the lowISRF (*left*), ref (*centre*) and hiZ (*right*) simulations. The white ellipse in each panel indicates the boundary of the cloud, defined by gas particles with a density above a threshold $n_{\text{H}, \text{min}} = 10 \text{ cm}^{-3}$. The black contours in the top row show $I_{\text{CO}} = 0.25 \text{ K km s}^{-1}$, which corresponds to the 3σ intensity threshold for the Small Magellanic Cloud in the observations of Leroy et al. (2011). In the examples from the ref and hiZ simulations, only the centres of the clouds would be detectable in a typical CO survey.

(see section 2.1). While the implementation of this pressure floor was motivated by numerical reasons, its effect on the cloud will be similar to a pressure term from unresolved turbulence. We can attribute a one-dimensional velocity dispersion, $\sigma_{\text{floor}, 1\text{D}}$, to the pressure floor according to $P_{\text{floor}} = \rho \sigma_{\text{floor}, 1\text{D}}^2$. Using equation 2.1 for P_{floor} , with our fiducial parameters $N_{\text{J}, \text{m}} = 4$, $m_{\text{gas}} = 750 \text{ M}_{\odot}$ and $N_{\text{ngb}}^{\text{SPH}} = 100$, we find:

$$\sigma_{\text{floor}, 1\text{D}} = 1.18 \left(\frac{\rho}{10^{-24} \text{ g cm}^{-3}} \right)^{1/6} \text{ km s}^{-1}. \quad (3.3)$$

We therefore include microturbulent broadening due to this pressure floor when computing the CO line emission, with a velocity dispersion given by equation 3.3. The Doppler broadening due to this microturbulence is then added to the thermal broadening in quadrature.

For each cloud in our simulations, we extracted a region around the cloud and interpolated the gas density, temperature and velocities, along with the densities of CO and H_2 , onto a 3d Cartesian grid with a resolution of 1 pc, using the same C^2 Wendland (1995) kernel with 100 SPH neighbours as was used in the simulations. We used RADMC-3D to compute the total emission from the $J = 1-0$ line and thermal dust emission in 480 wavelength bins covering a velocity range $\pm 60 \text{ km s}^{-1}$ centred on the line, which we projected onto a plane parallel to the galactic disc. We then repeated this without line emission to create a map of the thermal dust emission only, which we finally subtracted from the total emission to produce a continuum-subtracted map of CO $J = 1-0$ line emission.

In Fig. 5 we show examples of CO emission maps from individual molecular clouds in the lowISRF, ref and hiZ simulations (left, centre and right columns, respectively). The top and bottom rows show maps of the CO emission and gas surface density, respectively.

The white ellipse in each panel shows the boundary of the cloud, defined by gas particles with a density above a threshold $n_{\text{H}, \text{min}} = 10 \text{ cm}^{-3}$. This ellipse was computed by

projecting the 3d ellipsoid that approximately encloses the particles in the cloud onto the image plane, where the 3d ellipsoid is based on the principle axes of the moment of inertia tensor, as described in section 3.1. For comparison, the black contours in the top row show $I_{\text{CO}} = 0.25 \text{ K km s}^{-1}$. This corresponds to the 3σ intensity threshold for the observations of the Small Magellanic Cloud in Leroy et al. (2011). In the examples from the ref and hiZ simulations (centre and right columns, respectively), only the centres of the clouds are above this detection threshold.

We thus see that our standard definition of a cloud, based on a fixed density threshold, includes a larger region than if we had defined clouds based on the observable CO emission. We therefore also consider an alternative cloud definition, in which we only include regions in the 2d maps of CO emission with $I_{\text{CO}} > 0.25 \text{ K km s}^{-1}$. For this alternative cloud definition, we also compute the projected cloud mass, M_{proj} , and size, $R_{\text{proj}} = (A/\pi)^{1/2}$, from the 2d maps, rather than from the 3d particle distribution, where M_{proj} and A are the total mass and area, respectively, of pixels above the CO intensity threshold. This alternative cloud definition provides a fairer comparison with observations.

It is important to note that these CO emission maps may be sensitive to resolution. In particular, high-resolution simulations of dense clouds find that most CO is concentrated in compact structures, with sizes $\sim 1 \text{ pc}$ and densities $\sim 10^3 \text{ cm}^{-3}$ (e.g. Glover & Clark 2012). Such structures are poorly resolved in our simulations, even in our high resolution simulation in Appendix A, which may make the predicted CO emission uncertain.

4 CLOUD SCALING RELATIONS

Observations of molecular clouds find strong relations between their properties such as size, velocity dispersion and mass, both in Milky Way GMCs (e.g. Larson 1981; Solomon et al. 1987; Heyer et al. 2009) and in extragalactic GMCs (e.g. Bolatto et al. 2008). For example, building on the original relations identified by Larson (1981), Solomon et al. (1987) studied a sample of GMCs in the Milky Way, and found that the line of sight velocity dispersion, σ , follows a power law relation with cloud radius, R :

$$\sigma = 0.72 \left(\frac{R}{\text{pc}} \right)^{0.5} \text{ km s}^{-1}. \quad (4.1)$$

By assuming that the clouds are in virial equilibrium, they estimated the cloud masses, M , and they found a power-law relation between M and σ :

$$M = 2000 \left(\frac{\sigma}{\text{km s}^{-1}} \right)^4 \text{ M}_{\odot}. \quad (4.2)$$

Combining equations 4.1 and 4.2 produces the following relation between M and R :

$$M = 540 \left(\frac{R}{\text{pc}} \right)^2 \text{ M}_{\odot}, \quad (4.3)$$

which implies that the clouds in their sample have a constant mean surface density of $\Sigma = 170 \text{ M}_{\odot} \text{ pc}^{-2}$. Later studies have corrected this value to $200 \text{ M}_{\odot} \text{ pc}^{-2}$ to account for




Figure 6. Cloud mass, M_{cloud} , versus cloud radius, $R_{\text{mean}} = (abc)^{1/3}$, for all clouds with at least 50 particles identified in snapshots at 100 Myr intervals from 100 Myr to 900 Myr in the ref simulation. Our other two simulations (not shown) have very similar cloud mass-size relations. The colour scale indicates the cloud age. The black lines show the observed relations of Solomon et al. (1987) (solid; our equation 4.4) and Roman-Duval et al. (2010) (dashed; our equation 4.5). Our simulated clouds follow a similar slope to the observed relations, but the normalisation, which depends on the cloud definition, is a factor ≈ 4 lower than is observed.

an updated estimate for the Sun’s galactocentric radius of 8.5 kpc, rather than 10 kpc as originally used (see e.g. Heyer et al. 2009). The corrected mass-size relation is then:

$$M = 625 \left(\frac{R}{\text{pc}} \right)^2 M_{\odot}. \quad (4.4)$$

The updated galactocentric radius of the Sun will also affect the size-linewidth relation in equation 4.1. However, the correction for this relation is smaller than the errors.

Heyer et al. (2009) re-examined the sample of Solomon et al. (1987), and used the ^{13}CO luminosities of the clouds to estimate their molecular hydrogen masses. They found a median molecular surface density of $42 M_{\odot} \text{ pc}^{-2}$, lower than the value determined by Solomon et al. (1987) from the virial mass. Furthermore, they found that Σ is not constant, and that $\sigma/R^{0.5}$ varies systematically with surface density as $\Sigma^{0.5}$.

Roman-Duval et al. (2010) studied the properties of molecular clouds in the BU-FCRAO Galactic Ring Survey (Jackson et al. 2006) and the UMSB survey (Clemens et al. 1986; Sanders et al. 1986) in the Milky Way. They found the following mass-size relation:

$$M = 228 \left(\frac{R}{\text{pc}} \right)^{2.36} M_{\odot}, \quad (4.5)$$

based on ^{13}CO line emission.

Fig. 6 shows the cloud mass-size relation from the ref simulation, where the cloud radius, R_{mean} , was calculated from the 3d particle distribution (see equation 3.2). Our other two simulations (lowISRF and hiZ; not shown) have

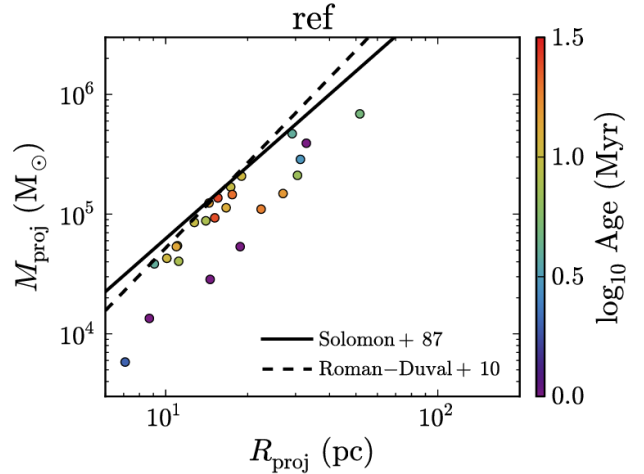


Figure 7. As Fig. 6, but for cloud masses and projected sizes determined from CO-detectable ($I_{\text{CO}} > 0.25 \text{ K km s}^{-1}$) regions only. The normalisation of this relation in our simulations is a factor ≈ 2 lower than the observed relations (Solomon et al. 1987 and Roman-Duval et al. 2010; black solid and dashed lines respectively) when we include only CO-detectable regions, compared to a factor of four lower than is observed when we used a density-based cloud definition (see Fig. 6).

very similar cloud mass-size relations. We show all clouds with at least 50 particles identified in snapshots at 100 Myr intervals from 100 Myr to 900 Myr. The colour scale indicates the age of the cloud, defined from the particles originally in the cloud in the current snapshot, as described in section 3.2. The black solid and dashed lines show the observed relations of Solomon et al. (1987) and Roman-Duval et al. (2010), respectively, i.e. equations 4.4 and 4.5.

Our simulated clouds follow a similar slope to the observed relations, but the normalisation is a factor ≈ 4 lower than is observed. The lower normalisation is determined by the density threshold that we use to define a cloud, and suggests that our definition includes a larger region around the cloud than would be included in a typical observational survey based on CO emission. Indeed, we saw in Fig. 5 that only the central regions of our simulated clouds have velocity-integrated CO intensities above a threshold of 0.25 K km s^{-1} . For a fairer comparison with observations, we therefore also considered an alternative cloud definition in which we include only regions above this CO intensity threshold, and compute cloud properties in projection, as described in section 3.3.

Fig. 7 shows the mass-size relation computed using only CO-detectable regions, for the ref simulation. Compared to Fig. 6, for a density-based cloud definition, the clouds lie much closer to the observed relations of Solomon et al. (1987) and Roman-Duval et al. (2010) (black solid and dashed lines, respectively), although the normalisation of this relation in our simulations is still a factor ≈ 2 lower than is observed (compared to a factor ≈ 4 in Fig. 6). However, even our CO-based cloud definition is not identical to the definitions used in these two observational studies. Our criterion is based on a minimum velocity-integrated CO intensity in the projected, two-dimensional position-




Figure 8. One-dimensional velocity dispersion, measured from gas particle motions only (σ_{1D} ; *top panel*), and including the contribution from the pressure floor ($\sigma_{\text{floor},1D}$; *bottom panel*), plotted against cloud radius. The colour scale in the top panel indicates the cloud age, while in the bottom panel it indicates the distance of the cloud’s centre of mass from the galaxy centre. The black lines indicate the observed relation of [Solomon et al. \(1987\)](#). We show only the ref simulation here, although the relations in our other two simulations are very similar. If we do not account for the pressure floor (top panel), we find unrealistically low velocity dispersions ($< 1 \text{ km s}^{-1}$). When we include the pressure floor as an unresolved turbulence term (bottom panel), we find velocity dispersions $> 2.3 \text{ km s}^{-1}$ for all clouds, although there remains a large scatter in this relation. The clouds with the highest velocity dispersion tend to be close to the galaxy centre.

position space of the CO emission maps. However, [Solomon et al. \(1987\)](#) define clouds as regions above a minimum CO brightness temperature of 1 K in the three-dimensional position-position-velocity (PPV) space. [Roman-Duval et al. \(2010\)](#) use a minimum velocity-integrated intensity of $4\sigma = 0.23\sqrt{N_\nu} \text{ K km s}^{-1}$, where N_ν is the number of velocity channels, but for ^{13}CO line emission, rather than ^{12}CO as used by us. Additionally, when measuring ^{13}CO column densities to compute the cloud mass, they include only velocity channels above a ^{13}CO brightness temperature of $4\sigma = 1 \text{ K}$. Therefore, the remaining discrepancy in the normalisation of the cloud mass-size relation is likely due to the different CO thresholds that we use. Given that observational studies use a range of cloud definitions, with different clump finding algorithms and using different molecular emission lines, we keep our CO-based definition general, rather than try to match a particular observational study.

[Pan et al. \(2016\)](#) used a hydrodynamic simulation of a barred spiral galaxy to investigate how the definition of




Figure 9. As Fig. 8, but measured from CO-detectable ($I_{\text{CO}} > 0.25 \text{ K km s}^{-1}$) regions of each cloud only. Also, we use projected sizes and compute the velocity dispersion by fitting a single Gaussian component to the CO spectrum. We exclude all clouds that show multiple peaks in their CO spectrum, as they cannot be fit with a single velocity component, and likely consist of multiple clouds that are undergoing mergers. We find better agreement with the observed relation of [Solomon et al. \(1987\)](#) (black line) than we saw in Fig. 8.

GMCs in position-position-position (PPP) or PPV space affects their properties. They found that the power law indices of the cloud scaling relations vary with cloud definition and, for a PPV-based definition, with disc inclination. [Duarte-Cabral & Dobbs \(2016\)](#) also explored different definitions of GMCs in PPP or PPV space, based on either H_2 or CO, using a high-resolution ($3.85 M_\odot$ per SPH particle) re-simulation of a section of a spiral galaxy simulation. They found that a PPV-based definition tends to blend clouds that would be physically separated in PPP space, although this effect is less significant if the PPV emission maps have high spatial resolution and high sensitivity. They also found that CO densities tend to trace only the high-density H_2 gas, rather than all molecular gas.

Fig. 8 compares the velocity dispersion-cloud size relation from the ref simulation (coloured points) to the observed relation from [Solomon et al. \(1987\)](#), i.e. our equation 4.1 (black solid line), for our standard density-based cloud definition. In the top panel of Fig. 8, we plot the one-dimensional velocity dispersion, $\sigma_{1D} = \sigma_{3D}/\sqrt{3}$, where the three-dimensional velocity dispersion is $\sigma_{3D}^2 = \sigma_x^2 + \sigma_y^2 + \sigma_z^2$, and $\sigma_i^2 = \langle v_i^2 \rangle - \langle v_i \rangle^2$ for the i^{th} component of the particle velocities, v_i , where angular brackets indicate a mass-weighted average over all particles in the cloud. The colour scale in the top panel of Fig. 8 indicates the cloud age.

We see that the relation between σ_{1D} and R_{mean} is steeper than observed. In particular, we find very low velocity dispersions, $\sigma_{1D} < 1 \text{ km s}^{-1}$, far below the observed relation of [Solomon et al. \(1987\)](#). However, these measurements of σ_{1D} in the simulations do not account for unresolved turbulence. As noted in section 3.3, the pressure floor that we impose on the gas to ensure that the Jeans mass is always well-resolved will have a similar effect on the cloud as

a pressure term from unresolved turbulence, with a turbulent velocity dispersion, $\sigma_{\text{floor}, 1\text{D}}$, given by equation 3.3. We therefore need to include the effects of this pressure floor.

In the bottom panel of Fig. 8, we compute $\sigma_{\text{floor}, 1\text{D}}$ for each cloud using its mean density, add this to $\sigma_{1\text{D}}$ in quadrature, and plot the total velocity dispersion against cloud radius. By accounting for the pressure floor in this way, we avoid unrealistically low velocity dispersions (the lowest value is now $\approx 2.3 \text{ km s}^{-1}$). This will be important for computing CO emission in our simulated clouds, as the CO $J = 1-0$ line is often optically thick in molecular clouds, so the line intensity will depend on the line width.

The observed velocity dispersions will also include a component due to the thermal broadening of the molecular lines that are used to measure $\sigma_{1\text{D}}$. For CO, with a mean molecular weight $\mu = 28$, the thermal velocity is $\sigma_{\text{th}, 1\text{D}} = \sqrt{k_{\text{B}}T/\mu m_{\text{p}}} = 0.17 \text{ km s}^{-1}$ at a temperature $T = 100 \text{ K}$, where k_{B} and m_{p} are the Boltzmann constant and proton mass, respectively. Thus, $\sigma_{\text{th}, 1\text{D}}$ is small compared to $\sigma_{\text{floor}, 1\text{D}}$ in our simulations, so we do not include $\sigma_{\text{th}, 1\text{D}}$ in Fig. 8, although we do account for thermal broadening when we compute CO line emission, as described in section 3.3.

Even accounting for the pressure floor, we still find a lot of scatter in this relation in our simulations, with some clouds showing velocity dispersions $> 10 \text{ km s}^{-1}$. In the top panel of Fig. 8, we found no trend in this relation with the cloud age. However, in the bottom panel, the colour scale indicates the distance of the cloud's centre of mass from the centre of the galaxy. We see that clouds with the highest velocity dispersions ($\gtrsim 10 \text{ km s}^{-1}$) are generally found within the central $\approx 1 \text{ kpc}$ of the galaxy. We also find that many of these high velocity dispersion clouds contain multiple density peaks that indicate substructures within the cloud. Therefore, some of the scatter towards high velocity dispersions is likely to be caused by motions of substructures within the cloud, possibly created by ongoing cloud-cloud mergers, which are more common in the centre of the galaxy. Interestingly, observations of molecular clouds in the centre of the Milky Way also find higher velocity widths compared to the linewidth-size relation of nearby molecular clouds in the Galactic disc (e.g. Oka et al. 2001).

Fig. 9 shows the velocity dispersion-size relation in the ref simulation using our CO-based cloud definition, i.e. restricted to CO-detectable regions and with cloud sizes computed in projection. The 1d velocity dispersion of each cloud was measured by fitting a single Gaussian component to the CO spectrum extracted from pixels above the I_{CO} threshold. We visually inspected each spectrum and excluded those with multiple peaks, which cannot be fit with a single velocity component. These systems are likely multiple clouds that are undergoing mergers. The velocity dispersion was then obtained from the width of the best-fitting Gaussian. The width of the CO spectrum includes microturbulent Doppler broadening by the pressure floor. By defining clouds above a CO intensity threshold, measuring the velocity dispersion from the width of the CO spectrum rather than motions of the gas particles, and excluding merging systems with multiple velocity components, we find better agreement with the observed relation of Solomon et al. (1987) than we saw in Fig. 8.

In Fig. 10 we plot the virial parameter, α , against M_{cloud} , from the ref simulation, where:

$$\alpha = \frac{5\sigma^2 R_{\text{mean}}}{GM_{\text{cloud}}} \quad (4.6)$$

(e.g. Bertoldi & McKee 1992; Dobbs et al. 2011), and we include the pressure floor in the velocity dispersion, i.e. $\sigma^2 = \sigma_{1\text{D}}^2 + \sigma_{\text{floor}, 1\text{D}}^2$. The numerical factor on the right hand side depends weakly on the density profile of the cloud. The value of 5 that we use here corresponds to a cloud with constant density; for comparison, in a cloud with a power-law density profile $\rho \propto r^{-2}$, this numerical factor would be 3.

The horizontal dotted line shows $\alpha = 1$, which corresponds to virial equilibrium, with $2K + W = 0$, where K and W are the kinetic and gravitational potential energies, respectively. A cloud that is gravitationally bound, with $K + W < 0$, requires $\alpha < 2$. While we do find clouds with $\alpha \approx 1 - 2$ in our simulations, which are marginally bound (but not virialised), most have $\alpha > 2$, and thus are unbound. Dobbs et al. (2011) similarly found that most (but not all) GMCs in their simulations of isolated disc galaxies are gravitationally unbound. They attributed this to cloud-cloud collisions and stellar feedback, which regulate the velocity dispersion within the clouds. However, in our simulations, the lack of clouds with low virial parameters is partially due to the pressure floor, at least for masses $\lesssim 3 \times 10^5 M_{\odot}$. We find a lower envelope of $\alpha \propto M^{-2/3}$ in Fig. 10, whereas observations find that α is approximately constant with mass (e.g. Rosolowsky 2007). This scaling of α with cloud mass is what we would expect when the virial parameter is determined by the pressure floor, with $\sigma_{\text{floor}, 1\text{D}} \propto \rho^{1/6}$ (equation 3.3) and $R_{\text{mean}} \propto M^{1/2}$ (as seen in Fig. 6). It is therefore apparent that the pressure floor prevents the low-mass clouds from becoming gravitationally bound in our simulations.

Some observational studies also suggest that molecular clouds may be gravitationally unbound (e.g. Heyer et al. 2001). Dobbs et al. (2011) also demonstrated that many of the GMCs in the sample of Heyer et al. (2009) have $\alpha > 2$ (see, for example, the centre bottom panel of fig. 1 of Dobbs et al. 2011). However, there are still some GMCs in this sample with $\alpha < 1$, which we do not see in our simulations. Furthermore, other studies suggest that molecular clouds may be marginally gravitationally bound, with $\alpha \approx 1$ (see e.g. McKee & Ostriker 2007).

To test how the pressure floor affects our results, we repeated the ref model twice, with the pressure floor lowered by factors of 4 and 16 in terms of the Jeans mass, corresponding to $N_{\text{J}, \text{m}} = 1$ and 0.25, respectively. We present these comparisons in Appendix B, and we also present resolution tests, with the mass resolution increased/decreased by a factor of four, in Appendix A. We summarise the main results here.

As we lower the pressure floor, the low-mass ($\lesssim 3 \times 10^5 M_{\odot}$), most poorly resolved clouds become more compact. They extend to lower values of α and can become strongly gravitationally bound, with $\alpha < 1$, and thus they can survive for longer, with ages up to $\approx 50 \text{ Myr}$. However, clouds with higher masses than this are unaffected by the pressure floor. This also means that the mass-size relation becomes flatter when we lower the pressure floor, and no longer agrees with the observed slope of this relation. Fur-




Figure 10. Virial parameter, α , versus cloud mass for all clouds with more than 50 particles in snapshots at 100 Myr intervals from 100 Myr to 900 Myr, from the ref simulation. Clouds in our other simulations (lowISRF and hiZ; not shown) have similar distributions of α . The colour scale indicates the cloud age. The horizontal dotted line shows $\alpha = 1$, which corresponds to virial equilibrium. We see that all of our simulated clouds have $\alpha > 1$, i.e. they are not virialised. However, the lower envelope of points is caused by the pressure floor, which prevents low-mass clouds ($\lesssim 3 \times 10^5 M_\odot$) from becoming strongly gravitationally bound.

thermore, these trends are not seen in our resolution tests. In particular, our highest resolution run reproduces the observed slope of the mass-size relation, and there are no clouds with $\alpha < 1$. Therefore, it is likely that the compact, long-lived clouds with $\alpha < 1$ that we find when we lower the pressure floor are artifacts of spurious fragmentation and collapse that may arise when we do not fully resolve the Jeans scale (e.g. Bate & Burkert 1997; Truelove et al. 1997).

Despite the differences that arise from lowering the pressure floor, we find that the median relations of molecule abundances with cloud age, and of CO intensity and X_{CO} factor with dust extinction, which we present for our fiducial simulations in the next two sections, are insensitive to the pressure floor, although the scatter in these relations does increase as we lower the pressure floor. These relations are also insensitive to resolution, although the scatter is higher at higher resolution.

Fig. 11 shows the virial parameter plotted against cloud mass for our CO-based cloud definition, including only regions above the I_{CO} threshold, and using velocity dispersions measured by fitting a single Gaussian component to the simulated CO spectra, as described above. We show clouds from the ref simulation (using our fiducial pressure floor, with $N_{\text{J},m} = 4$), and we exclude those with multiple peaks in the CO spectrum, which cannot be fit by a single Gaussian component. Compared to Fig. 10 (for a density-based cloud definition), the dependence of the lower envelope of α on cloud mass is weaker, which suggests that the impact of the pressure floor on the virial parameter is less severe when we use a CO-based cloud definition. However, we still find that all clouds in our simulations are unvirialised, with $\alpha > 1$.




Figure 11. As Fig. 10, but measured from CO-detectable regions only, and with α calculated from velocity dispersions measured from the simulated CO spectra. Compared to Fig. 10, for a density-based cloud definition, the dependence of the lower envelope of α on cloud mass is weaker. However, we still find that all clouds are unvirialised ($\alpha > 1$), even for a CO-based cloud definition.

5 CHEMICAL EVOLUTION

We now look at the evolution of molecular abundances within the dense clouds that we have identified in our simulations. In particular, we investigate the time-scales over which these clouds become fully molecular, and whether they remain close to chemical equilibrium throughout their evolution. As in the previous section, we consider two cloud definitions: one based on a minimum density threshold ($n_{\text{H}} > 10 \text{ cm}^{-3}$), and one based on a minimum velocity-integrated CO intensity threshold ($I_{\text{CO}} > 0.25 \text{ K km s}^{-1}$). We consider two molecular species: H_2 , which is the most prevalent molecule in interstellar gas, and CO, which is the most easily observed molecule.

5.1 Molecular hydrogen

Fig. 12 shows the molecular hydrogen fraction, $f_{\text{H}_2} = M_{\text{H}_2}/M_{\text{H,tot}}$ (where M_{H_2} is the mass of H_2 and $M_{\text{H,tot}}$ is the total mass of hydrogen in the cloud), for all clouds with at least 50 particles identified in snapshots at 100 Myr intervals from 100 Myr to 900 Myr, using our density-based cloud definition. In the top row of Fig. 12 we plot f_{H_2} against the age of the cloud, and in the bottom row we plot the ratio $f_{\text{H}_2}/f_{\text{H}_2}^{\text{eqm}}$, where $f_{\text{H}_2}^{\text{eqm}}$ is the molecular hydrogen fraction if all gas particles are set to chemical equilibrium. The three columns in Fig. 12 correspond to our three simulations (lowISRF, ref and hiZ), and the colour of each point indicates the mass of the cloud. The black curve in each panel shows the median in bins of age.

The top row of Fig. 12 shows that the H_2 fraction increases with age, while there does not appear to be any significant trend with cloud mass. The simulation using solar metallicity (hiZ, right column) shows the highest H_2 fractions for a given cloud age. This is as expected, as we as-




Figure 12. Molecular hydrogen fraction, f_{H_2} (top row), and the ratio of the H_2 fraction to the H_2 fraction in chemical equilibrium, $f_{\text{H}_2}/f_{\text{H}_2}^{\text{eqm}}$ (bottom row), plotted against cloud age for all clouds with at least 50 particles identified in snapshots at 100 Myr intervals from 100 Myr to 900 Myr in our three simulations: lowISRF (left column), ref (centre column) and hiZ (right column). The dotted horizontal lines indicate a value of unity, and the colour scale indicates cloud mass. We also show the median f_{H_2} or $f_{\text{H}_2}/f_{\text{H}_2}^{\text{eqm}}$ in bins of age (solid curves). We see that f_{H_2} increases, and moves closer to chemical equilibrium, with increasing cloud age, and does so faster for higher metallicity.

sume that the formation rate of H_2 on dust grains scales linearly with metallicity. In contrast, our reference simulation (ref, centre column), with a metallicity of ten per cent solar, shows the lowest H_2 fractions.

In run hiZ, the median cloud H_2 fraction reaches $f_{\text{H}_2} = 0.5$ after ≈ 3 Myr. In the ref simulation, with a factor of ten lower metallicity than hiZ, molecular hydrogen takes longer to build up, and only reaches $f_{\text{H}_2} \approx 0.2$ after ≈ 30 Myr. Finally, in run lowISRF (left column), with ten per cent solar metallicity and ten per cent of the ISRF used in the other two simulations, the median H_2 fraction is always a factor ≈ 2 higher than for ref. Thus, the time-scale for forming molecular H_2 in dense clouds is shorter at higher metallicity and (to a lesser extent) in the presence of a weaker UV radiation field.

In the bottom row of Fig. 12, we see that the H_2 fraction in young clouds is below what we would expect in chemical equilibrium. The clouds in the run at solar metallicity (hiZ) reach chemical equilibrium the fastest, with the median $f_{\text{H}_2}/f_{\text{H}_2}^{\text{eqm}}$ already at 50 per cent after ≈ 1 Myr (which is the smallest time-scale that we show here, as we only have snapshots at 1 Myr intervals). After ≈ 13 Myr, f_{H_2} has reached 90 per cent of its equilibrium value.

At lower metallicity, clouds take longer to reach chemical equilibrium. For example, clouds in the ref and lowISRF simulations reach 50 per cent of the equilibrium H_2 fraction after ≈ 16 Myr, and they reach 90 per cent after

≈ 22 Myr and ≈ 30 Myr respectively. In the ref simulation, clouds still have a low H_2 fraction ($f_{\text{H}_2} \approx 0.2$) after 30 Myr, although they have reached chemical equilibrium by this time. In other words, these clouds are still not fully molecular, even in chemical equilibrium. This suggests that, in the reference simulation, the H I -to- H_2 transition, which depends on both metallicity and radiation field, lies further above the density threshold that we use to define our clouds ($n_{\text{H}, \text{min}} = 10 \text{ cm}^{-3}$) than in the other two simulations. In the ref simulation our definition of a dense cloud therefore includes a greater proportion of the H I envelope.

In Fig. 13 we repeat Fig. 12, but for our CO-based cloud definition, i.e. including only regions with $I_{\text{CO}} > 0.25 \text{ K km s}^{-1}$. We compute f_{H_2} by projecting the H_2 column density onto the same image grid as was used for the CO emission maps, and selecting pixels above the I_{CO} threshold.

In the top row of Fig. 13, the H_2 fraction in the lowISRF and ref simulations shows less scatter than we previously saw in Fig. 12, and the values of f_{H_2} are higher, as we exclude the outer atomic envelope of the cloud. The ratio of H_2 fraction in non-equilibrium and H_2 fraction in equilibrium in the bottom row of Fig. 13 also shows less scatter.

The lowISRF run shows similar trends with cloud age as previously, whereas the ref run shows weaker evolution with age, with H_2 fractions closer to equilibrium (within a factor $\approx 2 - 5$) in young clouds when we include only CO-detectable regions. The ref simulation contained the most




Figure 13. As Fig. 12, but only for the region within each cloud with $I_{\text{CO}} > 0.25 \text{ K km s}^{-1}$. By restricting the cloud definition to CO-detectable regions, we find less scatter in the H_2 fraction for the lowISRF and ref simulations, while the hiZ simulation is mostly unaffected.

CO-faint pixels, because its combination of low metallicity and high radiation field resulted in the lowest CO fractions (see Fig. 14 in the next section). Therefore, restricting our cloud definition to CO-detectable regions has the strongest effect for the ref run. The CO-detectable regions are located in the dense cores of the clouds, which we would expect to reach chemical equilibrium faster, since collisional reaction rates typically scale with n^2 , where n is the density. This likely explains why the H_2 fraction in the ref simulation reaches equilibrium faster when we select only CO-detectable regions.

In the hiZ simulation, the H_2 fraction is mostly unaffected by restricting the cloud definition to CO-detectable regions, with similar scatter and trends with age as in Fig. 12.

One caveat to note is that these conclusions on the formation time-scale of H_2 may be sensitive to resolution. In particular, small-scale turbulence makes the gas form dense clumps, which increases the formation rate of H_2 in turbulent clouds (e.g. Glover & Mac Low 2007b; Micic et al. 2012). However, our simulations do not resolve this small-scale turbulence, even in the high resolution test in Appendix A, so it is likely that we underestimate the formation rate of H_2 .

Krumholz & Gnedin (2011) compared the equilibrium H_2 model of Krumholz et al. (2008, 2009) and McKee & Krumholz (2010) to the non-equilibrium H_2 model of Gnedin & Kravtsov (2011), applied to cosmological zoom-in simulations of a Milky Way progenitor galaxy and to a simulation of a cosmological box, $25h^{-1} \text{ Mpc}$ on a side, run with the Adaptive Refinement Tree (ART) code (Kravtsov 2003).

They found excellent agreement at metallicities $\gtrsim 10^{-2} Z_{\odot}$, suggesting that non-equilibrium effects are unimportant for H_2 at the metallicities that we consider here. However, their simulations were run at a lower resolution than we use. For example, their cosmological zoom-in simulations had a maximum resolution of 65 pc, compared to a gravitational softening of 3.1 pc in our simulations.

Hu et al. (2016) explored the effects of non-equilibrium chemistry on the H_2 fraction in hydrodynamic simulations of dwarf ($M_{200} = 2 \times 10^{10} M_{\odot}$) galaxies at a much higher resolution than we use ($4 M_{\odot}$ per SPH particle). They find that the H_2 mass is out of equilibrium throughout their simulations, which agrees with our results.

5.2 Carbon monoxide

The top row of Fig. 14 shows the mass fraction of carbon in CO, $f_{\text{CO}} = \frac{12}{28} \frac{M_{\text{CO}}}{M_{\text{C, tot}}}$, for each cloud as a function of cloud age, where M_{CO} and $M_{\text{C, tot}}$ are the CO and total carbon masses respectively, while the bottom row shows $f_{\text{CO}}/f_{\text{CO}}^{\text{eqm}}$, where $f_{\text{CO}}^{\text{eqm}}$ is the CO mass fraction in chemical equilibrium. The colour scale indicates cloud mass, the black curves show the median in bins of age, and the left, centre and right columns show runs lowISRF, ref and hiZ respectively.

In runs lowISRF and ref, which both assume $0.1 Z_{\odot}$, we see that f_{CO} tends to increase with cloud age. However, there is more scatter in f_{CO} at fixed age than we saw for f_{H_2} in Fig. 12. A handful of clouds reach $f_{\text{CO}} \approx 0.5$ in the lowISRF run, but many are several orders of magnitude below unity.




Figure 14. As Fig. 12, but for the CO fraction, $f_{\text{CO}} = (12/28)M_{\text{CO}}/M_{\text{C,tot}}$ (top row), and the ratio of the CO fraction to the equilibrium value, $f_{\text{CO}}/f_{\text{CO}}^{\text{eqm}}$. We define f_{CO} as the fraction, by mass, of carbon in CO molecules (hence the factor of $(12/28)$). The lowISRF and ref runs show increasing f_{CO} and $f_{\text{CO}}/f_{\text{CO}}^{\text{eqm}}$ with cloud age and mass, while the hiZ run shows no strong trends. For lowISRF and ref the CO fraction in young clouds is typically lower than in equilibrium, whereas for hiZ it is typically higher than in equilibrium.

In the bottom left and bottom centre panels, we see that f_{CO} in young clouds ($\lesssim 10 - 15$ Myr) tends to be below equilibrium by 1–2 orders of magnitude in the lowISRF and ref runs, while clouds older than this are typically closer to equilibrium, although for ref the median $f_{\text{CO}}/f_{\text{CO}}^{\text{eqm}}$ is still only 0.2–0.5. However, these non-equilibrium effects do not fully explain the very low CO fractions that we find in the top row. These low values of f_{CO} are partly due to the density threshold, $n_{\text{H,min}} = 10 \text{ cm}^{-3}$, that we use to define a cloud. This is close to the density of the H_I-to-H₂ transition, which can occur once H₂ becomes self-shielded. However, CO forms once it becomes shielded from dissociating radiation by dust, which typically occurs at higher densities.

The lowISRF and ref runs also show trends of f_{CO} with cloud mass, with more massive clouds showing higher CO fractions that are closer to equilibrium. This is because massive clouds are more likely to contain higher density regions where dust shielding is sufficient to form CO.

The simulation using solar metallicity (hiZ; right column) has higher CO fractions than the ref simulation. This is due to the higher dust abundance at higher metallicity, and hence stronger dust shielding from dissociating radiation. We see no strong trend of f_{CO} with cloud age in the hiZ simulation. In the bottom right panel, we see that the CO fraction is either close to equilibrium or enhanced, by up to two orders of magnitude in some cases. The enhanced CO fractions that we see in the hiZ run are due to fluctuations in the dust extinction seen by individual particles

within the cloud. Particles with enhanced CO abundances had $A_{\text{v}} \gtrsim 1$ within the previous few Myr, but A_{v} has since declined. Since the photodissociation rate of CO decreases exponentially with A_{v} , a small decrease in A_{v} can produce a large increase in photodissociation rate. However, it takes a finite time for the CO to be destroyed, thus we see enhanced CO abundances. We see much less enhancement of CO at lower metallicity (lowISRF and ref) because, in these runs, A_{v} rarely exceeds unity, thus CO rarely becomes fully shielded from dissociating radiation.

Fig. 15 shows CO fractions of clouds using a CO-based cloud definition, i.e. including only regions above the I_{CO} threshold. The effects of limiting our cloud definition to CO-detectable regions on CO fractions are similar to the effects it had on H₂ fractions that we saw in the previous section. In the lowISRF and ref runs, CO fractions are higher and show less scatter. The trends with cloud age in the lowISRF run are similar to those for a density-based cloud definition, while the ref simulation shows weaker evolution and is close to equilibrium, even in young clouds. The CO fractions in the hiZ run are mostly unaffected by the choice of cloud definition, and hence for young clouds they remain strongly enhanced compared to equilibrium.

6 CO EMISSION AND THE X_{CO} FACTOR

Observations of molecular clouds often use CO emission as a tracer of molecular gas. The H₂ column density is then




Figure 15. As Fig. 14, but only for the region within each cloud with $I_{\text{CO}} > 0.25 \text{ K km s}^{-1}$. The CO fraction in CO-detectable regions in the ref simulation is now close to equilibrium, even in young clouds, whereas with our previous cloud definition it was an order of magnitude below equilibrium in clouds younger than $\approx 10 \text{ Myr}$. However, for lowISRF and hiZ the median CO fraction in young clouds remains an order of magnitude lower and higher than in equilibrium, respectively.

determined from the CO intensity using a conversion factor, X_{CO} , as given in equation 1.1 (see Bolatto et al. 2013 for a recent review). If the abundances of H_2 and CO are out of equilibrium in young clouds, as we found to be the case in the previous section, then this may affect the X_{CO} factor. To investigate this, we used our maps of CO emission from the clouds in our simulations to measure X_{CO} .

The CO properties of a cloud are expected to depend on the dust extinction, as dust shields the cloud from photodissociating radiation, enabling the formation of CO (e.g. Lombardi et al. 2006; Pineda et al. 2008; Feldmann et al. 2012; Lee et al. 2015). Fig. 16 shows the mean velocity-integrated CO intensity (I_{CO} ; top row) and the mean X_{CO} factor (bottom row) in each cloud (using our density-based cloud definition) from our three simulations (lowISRF, ref and hiZ, in the left, centre and right columns respectively), as a function of the mean dust extinction, A_v . In each cloud, we average I_{CO} , A_v and the H_2 column density, N_{H_2} , over all pixels within the projected ellipse containing the cloud particles, i.e. the white ellipses in Fig. 5. The mean X_{CO} factor is then $\langle N_{\text{H}_2} \rangle / \langle I_{\text{CO}} \rangle$. The horizontal dotted line in the top row indicates $I_{\text{CO}} = 0.25 \text{ K km s}^{-1}$, which corresponds to the 3σ intensity threshold for the Small Magellanic Cloud in the observations of Leroy et al. (2011), and is the minimum I_{CO} threshold that we use in our CO-based cloud definition. In the bottom row, the horizontal dotted line shows a value of $X_{\text{CO}} = 2 \times 10^{20} \text{ cm}^{-2} (\text{K km s}^{-1})^{-1}$, typical of molecular

clouds in the Milky Way (Bolatto et al. 2013). The colour scale in both rows indicates cloud age.

In the top row, we see that I_{CO} increases steeply with A_v , particularly in the lowISRF and ref simulations. For comparison, we also show the observed $I_{\text{CO}} - A_v$ relations seen in the Pipe nebula (Lombardi et al. 2006) and the Perseus cloud (Pineda et al. 2008) in the Milky Way. The observations of Pineda et al. (2008) in particular find that I_{CO} cuts off at low A_v , below a threshold $A_v = 0.58$. This is unsurprising, as CO typically relies on dust to become shielded from dissociating radiation before it can form. Therefore, the steep $I_{\text{CO}} - A_v$ relation that we find in our simulations is likely due to this threshold effect, with most clouds lying close to the threshold. Since A_v depends on the column density, along with metallicity, this strong $I_{\text{CO}} - A_v$ relation reflects the fact that it is the column density, rather than the volume density, of a cloud that determines the molecular properties of the cloud, as this determines whether the cloud is shielded from dissociating radiation.

At high A_v , observations find that I_{CO} becomes saturated as the CO line becomes optically thick (e.g. Lombardi et al. 2006; Pineda et al. 2008). The lowISRF and ref simulations do not extend above $A_v \approx 0.4$ and so do not saturate, but the hiZ run contains clouds up to $A_v \approx 6$. These high- A_v clouds in the hiZ run do suggest a much shallower relation than at lower A_v in the same simulation, and are consistent with the observed saturation in Pineda et al. (2008), although we only have a few high- A_v clouds, so it is not




Figure 16. Mean velocity-integrated CO intensity, I_{CO} (top row), and X_{CO} factor (bottom row), plotted against mean dust extinction, A_v , for clouds from the lowISRF (left), ref (centre) and hiZ (right) simulations. We include all clouds with at least 50 particles in snapshots from 100 Myr to 900 Myr, in 100 Myr intervals. The colour scale indicates cloud age. In the top row, we also show the $I_{\text{CO}} - A_v$ relations observed in the Pipe nebula (Lombardi et al. 2006; solid curves) and the Perseus cloud (Pineda et al. 2008; dot-dashed curves) in the Milky Way. The horizontal dotted line in the top row indicates $I_{\text{CO}} = 0.25 \text{ K km s}^{-1}$, which corresponds to the 3σ intensity threshold for the Small Magellanic Cloud in the observations of Leroy et al. (2011). In the bottom row, the horizontal dotted line indicates the typical value measured in molecular clouds in the Milky Way, $X_{\text{CO}} = 2 \times 10^{20} \text{ cm}^{-2} (\text{K km s}^{-1})^{-1}$ (e.g. Bolatto et al. 2013).

clear if this relation is truly saturated in our simulations at high A_v .

Comparing the different panels in the top row, we see that the threshold A_v below which I_{CO} is strongly suppressed increases with metallicity (centre versus right) and, to a lesser extent, with the intensity of the radiation field (left versus centre). The dependence on radiation field is understandable, as a stronger radiation field requires a higher dust extinction before CO can become shielded.

However, the reason for the dependence on metallicity is more complicated. If CO is shielded only by dust, then the dissociation rate decreases $\propto \exp(-\gamma_d A_v)$, where $\gamma_d = 3.53$ (van Dishoeck et al. 2006). The threshold, A_v^{thresh} , then arises from the exponential cut off due to shielding. The formation of CO proceeds via a series of reactions, so the overall formation rate, R_{form} , will be determined by the rate-limiting step. These reactions are typically two-body interactions, so R_{form} scales with density squared. It also depends on the availability of carbon and oxygen, with densities $n_{\text{C,tot}}$ and $n_{\text{O,tot}}$ respectively, so $R_{\text{form}} \propto n_{\text{C,tot}} n_{\text{O,tot}} \propto Z^2 n_{\text{H,tot}}^2$, where Z is the metallicity and $n_{\text{H,tot}}$ is the total hydrogen number density. However, the rate-limiting step may depend on only Oxygen or Carbon, and not both (e.g. if the formation of an intermediate species such as CH_2^+ is the slowest step), in which case $R_{\text{form}} \propto Z n_{\text{H,tot}}^2$. If we define A_v^{thresh} to be when the CO fraction is some value, say

$f_{\text{CO}} = 0.1$, then $\exp(-\gamma_d A_v^{\text{thresh}}) \propto Z^i n_{\text{H,tot}}^2$, where $i = 1$ or 2 , depending on the rate-limiting step in the formation of CO. Since the clouds in all three of our simulations follow the same mass-size relation, and have the same distribution of cloud masses (Fig. 2), the average cloud surface density and volume density are independent of Z . Thus, the metallicity dependence of A_v^{thresh} is given by:

$$A_v^{\text{thresh}} = -\frac{i}{\gamma_d} \ln(Z) + \text{constant}. \quad (6.1)$$

We therefore expect A_v^{thresh} to decrease weakly with increasing metallicity, if the attenuation of CO photodissociation is due to dust shielding. However, this is opposite to what is seen in Fig. 16.

The reason for this discrepancy is that, in the ref simulation, with ten per cent solar metallicity, the shielding of CO is primarily due to H_2 , and not dust. H_2 shielding will cut off the CO dissociation rate at a threshold H_2 column density, $N_{\text{H}_2}^{\text{thresh}} = f_{\text{H}_2} N_{\text{H,tot}}^{\text{thresh}}$, where f_{H_2} is the H_2 fraction of the cloud and $N_{\text{H,tot}}^{\text{thresh}}$ is the total hydrogen column density at the threshold. Then $A_v^{\text{thresh}} \propto Z N_{\text{H,tot}}^{\text{thresh}}$. If $N_{\text{H,tot}}^{\text{thresh}}$ is constant, A_v^{thresh} will increase linearly with Z . However, $N_{\text{H}_2}^{\text{thresh}}$ decreases logarithmically with Z due to the increased CO formation rate, as described above, and f_{H_2} is generally higher in the hiZ run (Fig. 14). Addition-




Figure 17. As Fig. 16, but for I_{CO} , X_{CO} and A_V averaged over only pixels with $I_{\text{CO}} > 0.25 \text{ K km s}^{-1}$. Note that the ranges of the y-axes are much smaller than in Fig. 16. There is much less scatter in I_{CO} and X_{CO} at fixed A_V compared to Fig. 16, as we exclude CO-faint pixels that would be undetectable in typical CO surveys of molecular clouds.

ally, dust shielding becomes more important at high metallicity, which further reduces A_V^{thresh} , and CO self-shielding also plays a role in some clouds. We thus find a sub-linear increase in A_V^{thresh} with Z .

Pineda et al. (2008) find that the $I_{\text{CO}} - A_V$ relation in separate regions of the Perseus cloud also varies, suggesting that this relation depends on the physical conditions in the cloud. Using the Meudon PDR code (Le Petit 2006), they find that the variations in the $I_{\text{CO}} - A_V$ relation that they observe can be explained by variations in physical conditions, particularly volume density and internal gas motions. They also find that the $I_{\text{CO}} - A_V$ relation moves to higher A_V in the presence of stronger radiation fields in their models, consistent with what we see in our simulations. However, they do not consider variations in metallicity, which we find to be more important. Lee et al. (2015) measure the $I_{\text{CO}} - A_V$ relation in the Large and Small Magellanic Clouds, and compare these to the Milky Way. They find that, at fixed A_V , I_{CO} decreases with increasing dust temperature, suggesting a dependence on radiation field strength that is consistent with our simulations. However, they find that the $I_{\text{CO}} - A_V$ relations in these three galaxies are similar, despite their different metallicities.

The bottom row of Fig. 16 shows a large range in X_{CO} , spanning from two (lowISRF) to four (ref) orders of magnitude. We find no strong trends of X_{CO} with A_V . However, if we look at the highest- A_V clouds in the hiZ run ($A_V > 4$), the scatter in X_{CO} is much smaller, and the clouds lie within a factor two of the Milky Way value. We see a similar trend at

high A_V when we lower the pressure floor in the ref run (see the bottom right panel of Fig. B6). As we discuss further in Appendix B, when we lower the pressure floor the clouds become more compact and extend to higher A_V . In the run with the lowest pressure floor ($N_{\text{J,m}} = 0.25$), the scatter in X_{CO} at $A_V > 0.3$ is reduced by a factor four compared to the whole sample, and at $A_V > 0.6$ the clouds are consistent with the Milky Way value of X_{CO} . This suggests that the large scatter arises because the clouds are diffuse, with low A_V , and the scatter is greatly reduced at high A_V . However, this is based on a small number of clouds.

Bell et al. (2006) find a strong relation between X_{CO} and A_V in their one-dimensional PDR models. However, the $X_{\text{CO}} - A_V$ relations that they consider show how X_{CO} varies with depth in a given cloud, whereas in Fig. 16 we show the mean X_{CO} and A_V for individual clouds. Indeed, Bell et al. (2006) show that their $X_{\text{CO}} - A_V$ relation varies with cloud properties such as density and turbulent velocity. The scatter that we find in X_{CO} in our simulations is therefore likely driven by the wide range of cloud properties in our sample.

In the two runs at $0.1 Z_{\odot}$ (lowISRF and ref), we find a trend of increasing I_{CO} with cloud age. This is consistent with the trend of increasing f_{CO} with age that we saw in Fig. 14. These two runs also show a trend of decreasing X_{CO} with increasing cloud age. The median X_{CO} factor in bins of age decreases by more than an order of magnitude for 0 to 15 Myr, although there is a large scatter (two and four orders of magnitude in lowISRF and ref, respectively)

in X_{CO} at fixed age. The trend in X_{CO} at ages > 15 Myr is uncertain, as there are few clouds at high ages. We see no strong trends of I_{CO} or X_{CO} with age at solar metallicity (hiZ), as the time-scales to reach equilibrium are shorter at higher metallicity, as we saw in Figs. 12 and 14.

In Fig. 17 we plot I_{CO} (top row) and X_{CO} (bottom row) versus A_v for our CO-based cloud definition, where these three quantities are now averaged over only pixels within the cloud with $I_{\text{CO}} > 0.25 \text{ K km s}^{-1}$. Note that the ranges of the y-axes are much smaller than in Fig. 16. In the lowISRF and ref simulations (left and centre columns), there is much less scatter in both I_{CO} and X_{CO} than we saw in Fig. 16, for a density-based cloud definition. The reduction in scatter in the hiZ simulation is more modest. We also see more clearly how the $I_{\text{CO}} - A_v$ relation shifts towards higher A_v at higher metallicity and, to a lesser extent, at higher radiation field strength. While most high- A_v clouds in the hiZ run remain consistent with the observed saturation of I_{CO} , there are now two clouds that lie a factor of $\approx 2 - 3$ above the observed relation. We again find no strong trends of X_{CO} with A_v , except that clouds with $A_v > 6$ in the hiZ run show much less scatter and are within a factor two of the Milky Way value. In the lowISRF run, the trend of median X_{CO} with cloud age is similar to that using a density-based cloud definition, while the trend of X_{CO} with age in the ref run is much weaker when we use a CO-based cloud definition.

It is important to note that the values of the X_{CO} factor that we have presented in this section may be sensitive to resolution. In particular, as we noted in section 3.3, high-resolution simulations of dense clouds (e.g. Glover & Clark 2012) have found that most CO is concentrated in dense ($\sim 10^3 \text{ cm}^{-3}$), compact ($\sim 1 \text{ pc}$) structures that are poorly resolved in our simulations. This will make the predicted CO emission, and hence the X_{CO} factor, uncertain. Our resolution tests in Appendix A do not show any significant change in the $I_{\text{CO}} - A_v$ or $X_{\text{CO}} - A_v$ relations, except that there is more scatter at higher resolution. However, even our highest resolution run, with a gravitational softening of 2 pc, cannot resolve the $\sim 1 \text{ pc}$ structures that are likely to dominate the CO emission.

Glover & Clark (2016) studied the emission of CO and C I in high resolution ($0.005 M_\odot$ per SPH particle) simulations of individual molecular clouds with a range of metallicities, dust-to-gas ratios, UV radiation fields and cosmic ray ionisation rates. They found that, in low metallicity clouds ($Z \lesssim 0.2 Z_\odot$), X_{CO} decreases with time, while there is no strong evolution at high metallicity. This behaviour is similar to what we find in our simulations. They also demonstrated that C I is a better tracer of molecular gas than CO at low metallicity.

Smith et al. (2014) studied the molecular gas and CO properties in high resolution simulations of Milky Way-type galaxies, run with the moving mesh code AREPO (Springel 2010) and with time-dependent chemistry based on Glover & Mac Low (2007a,b) and Nelson & Langer (1997). They considered four simulations of galaxies with different surface densities and radiation field strengths, all assuming solar metallicity. The mean X_{CO} factor averaged over the galaxy in their four simulations was $3.89 - 13.1 \times 10^{20} \text{ cm}^{-2} (\text{K km s}^{-1})^{-1}$ when all gas was included. However, 42 – 85 per cent of the molecular mass in their simulations was ‘CO-dark’, with a CO intensity $\leq 0.1 \text{ K km s}^{-1}$.

When they included only gas with $I_{\text{CO}} > 0.1 \text{ K km s}^{-1}$, the mean X_{CO} factor in their simulations was $1.48 - 2.28 \times 10^{20} \text{ cm}^{-2} (\text{K km s}^{-1})^{-1}$. In other words, when they included only CO-bright regions, the mean X_{CO} factor decreased. This agrees with the trend that we find in our simulations. Furthermore, the values of the mean X_{CO} factor found by Smith et al. (2014), averaged over the whole galaxy, overlap with the lowest values that we find for individual clouds, although we also find a large scatter in X_{CO} between clouds.

7 CONCLUSIONS

We have presented an analysis of GMCs identified in high-resolution ($750 M_\odot$ per particle and a gravitational softening of 3.1 pc) SPH simulations of isolated, low-mass ($M_* \sim 10^9 M_\odot$) disc galaxies, with a particular emphasis on the evolution of molecular abundances and the implications for CO emission and the X_{CO} factor. Our simulations include a treatment for the non-equilibrium chemistry of 157 species, including 20 molecules (Richings et al. 2014a,b).

We define dense clouds in our simulations using two different methods, one that is physically motivated and another that is observationally motivated. First, we define clouds to be regions with a density above a threshold of $n_{\text{H}} = 10 \text{ cm}^{-3}$, which, depending on the metallicity and radiation field, is comparable to the density of the H I-to-H₂ transition (e.g. Schaye 2001; Gnedin et al. 2009). We group gas particles above this threshold into clouds using a Friends-of-Friends (FoF) algorithm, with a linking length of $l = 10 \text{ pc}$.

The observationally motivated cloud definition is based on CO emission, where we restrict the cloud boundary to regions with a velocity-integrated CO intensity above a threshold of 0.25 K km s^{-1} . This allows for a fairer comparison with observations, as it excludes CO-faint regions that would be undetectable in typical CO surveys of molecular clouds, although it is still not a like-for-like comparison.

To highlight the effects of metallicity and radiation field, we run our simulations at constant metallicity and with a uniform background UV radiation field, along with a prescription for local self-shielding by gas and dust (Richings et al. 2014b). Our three simulations cover two metallicities ($0.1 Z_\odot$ and Z_\odot) and two UV radiation fields (the ISRF of Black 1987, measured in the local solar neighbourhood, and ten per cent of this ISRF).

Our main results are as follows:

(i) Our simulated clouds have a median lifetime of 13 Myr (Fig. 4), defined by the period over which at least half of the particles originally in the cloud when it was identified were in the main progenitor/descendant. This is consistent with observational estimates (e.g. Bash et al. 1977; Kawamura et al. 2009; Murray 2011; Miura et al. 2012), although see Elmegreen (2000) and Scoville et al. (1979) for examples of shorter and longer observational estimates of GMC lifetimes, respectively. If we instead define the cloud lifetime by tracking the total mass of its main progenitor/descendant, rather than only the original particles, we find that clouds survive longer, with a median lifetime of 33 Myr, as new gas cycles through the cloud.

(ii) Simulated clouds follow a mass-size relation $M \propto R^2$, as observed in molecular clouds (e.g. Solomon et al. 1987; Roman-Duval et al. 2010). If we define clouds by a density

threshold of $n_{\text{H}} = 10 \text{ cm}^{-3}$, then the normalisation is a factor ≈ 4 below the observed relation (Fig. 6). However, if we restrict our cloud definition to CO-detectable regions, we find better agreement with observations, with a normalisation that is a factor ≈ 2 below the observed relation (Fig. 7).

(iii) Clouds defined by a density threshold approximately follow the observed velocity dispersion-size relation if we account for the contribution of the pressure floor in the velocity dispersion (Fig. 8), although there is a large scatter in the simulated relation. In particular, clouds within 1 kpc of the galactic centre typically lie 0.5 dex above the observed relation. We find better agreement with the observed relation when we use a CO-based cloud definition (Fig. 9).

(iv) Most clouds in our simulations are gravitationally unbound, with a virial parameter $\alpha > 2$ (Fig. 10). While some clouds are marginally bound, with $\alpha \approx 1 - 2$, all clouds have $\alpha > 1$, i.e. no clouds are virialised. This is partially due to the pressure floor, at least in low-mass clouds ($\lesssim 3 \times 10^5 M_{\odot}$). When we repeat the ref run with a lower pressure floor, reduced by a factor of 16 in terms of the Jeans mass, we do find some low-mass clouds with $\alpha < 1$ (Fig. B4). However, this may be an artifact of spurious fragmentation due to a poorly resolved Jeans scale, as we find no clouds with $\alpha < 1$ when we simultaneously increase the resolution and decrease the pressure floor (Fig. A4).

(v) At ten per cent solar metallicity, young GMCs ($\lesssim 10 - 15$ Myr) are underabundant in H_2 and CO compared to chemical equilibrium, by factors ≈ 3 and $1 - 2$ orders of magnitude respectively (Figs. 12 and 14). These non-equilibrium effects are less apparent at solar metallicity. The H_2 fraction at solar metallicity reaches within a factor 2 of equilibrium at 1 Myr, while the CO fraction at solar metallicity either remains close to equilibrium or becomes enhanced by up to two orders of magnitude compared to equilibrium. These non-equilibrium effects therefore depend strongly on metallicity, with no strong dependence on radiation field.

(vi) If we restrict our analysis to CO-detectable regions (with $I_{\text{CO}} > 0.25 \text{ K km s}^{-1}$), we find higher H_2 and CO fractions, as we exclude the atomic outer envelope of clouds (Figs. 13 and 15). The simulation with a low metallicity and a low UV radiation field (lowISRF) shows similar trends of f_{H_2} and f_{CO} with age as for our standard cloud definition, although the simulation at low metallicity and for a high UV radiation field (ref) shows weaker evolution of f_{H_2} and f_{CO} with age.

(vii) The mean CO intensity, I_{CO} , is strongly suppressed towards low dust extinction, A_{v} , and may become saturated at high A_{v} (Fig. 16), in agreement with observations (e.g. Pineda et al. 2008). Our simulated $I_{\text{CO}} - A_{\text{v}}$ relation moves towards higher A_{v} at higher metallicities and, to a lesser extent, for stronger UV radiation fields.

(viii) There is large scatter (2–4 orders of magnitude) in the mean X_{CO} factor of individual clouds (Fig. 16). Clouds at high A_{v} show much less scatter in X_{CO} and are within a factor of 2 of the Milky Way value (see also Fig. B6).

(ix) At ten per cent solar metallicity, we find weaker CO emission in young clouds, with ages $\lesssim 10 - 15$ Myr (Fig. 16), consistent with the trends we find for f_{CO} . This is also reflected in the median X_{CO} factor in bins of cloud age, which decreases by more than an order of magnitude from 0 to 15 Myr, although there is a large scatter in X_{CO} at fixed

age. There are no strong trends with age at solar metallicity.

(x) By restricting our analysis to CO-detectable regions, we find less scatter in X_{CO} ($\approx 1 - 2$ orders of magnitude; Fig. 17). We also find better agreement with observed GMC scaling relations (Figs. 7 and 9).

We have therefore shown that, at ten per cent solar metallicity, clouds younger than $\approx 10 - 15$ Myr are likely to be highly underabundant in H_2 and CO compared to chemical equilibrium. CO is more underabundant than H_2 in young clouds, which results in a trend of decreasing X_{CO} with increasing age from 0 to 15 Myr, albeit with a large scatter. Clouds at solar metallicity reach chemical equilibrium faster (within ≈ 1 Myr).

However, there are several caveats that we need to consider. Firstly, our simulations use a constant, uniform UV background. While we do include self-shielding by gas and dust, which varies both spatially and temporally, the interstellar radiation field itself should also vary, due to the birth and evolution of nearby stars and due to relative motions between the gas and the stars. If the fluctuations in the radiation field occur on time-scales shorter than the chemical time-scale, they can drive abundances out of equilibrium (e.g. Oppenheimer & Schaye 2013). Additionally, radiative feedback from young, massive stars (due to photoionisation heating and/or radiation pressure) may disrupt their natal clouds and inhibit further formation of dense clouds (e.g. Dale et al. 2005; Krumholz & Matzner 2009; Hopkins et al. 2012; Walch et al. 2012; Rosdahl et al. 2015). These effects are not included in our simulations.

Secondly, the presence of turbulence can drive chemical abundances out of equilibrium, if the time-scale associated with the turbulence is short compared to the chemical time-scale (e.g. Gray et al. 2015). If we do not fully resolve small-scale turbulence in our simulations, we may therefore underestimate the effects of non-equilibrium chemistry. Conversely, the presence of unresolved turbulence will also create small-scale regions with higher densities than we resolve. The formation time-scales of molecules are shorter at higher densities, so they would reach equilibrium sooner in these unresolved high-density regions. This could lead us to overestimate the non-equilibrium chemistry.

Furthermore, our simulations include a density-dependent pressure floor to ensure that we always resolve the Jeans mass by at least four times the kernel mass. However, in the lowest-mass clouds ($\lesssim 3 \times 10^5 M_{\odot}$) this pressure floor prevents clouds from becoming strongly gravitationally bound ($\alpha < 1$). In Appendix B we repeat our ref model with different pressure floors, to explore how this affects our results. As we lower the pressure floor, the lowest-mass, most poorly resolved clouds ($\lesssim 2 - 3 \times 10^5 M_{\odot}$) become more compact, more strongly gravitationally bound, and can live longer (up to ≈ 50 Myr). However, higher-mass clouds are unaffected. This also means that the run with the lowest pressure floor no longer reproduces the observed slope of the mass-size relation. This is inconsistent with what we find when we increase the resolution (Appendix A). Therefore, it is likely that the increasingly compact and gravitationally bound low-mass clouds that we find when we lower the pressure floor are the result of artificial fragmentation and

collapse that may arise when we do not fully resolve the Jeans scale (e.g. [Bate & Burkert 1997](#); [Truelove et al. 1997](#)).

Despite these differences when we vary the pressure floor, we find that our main results for the median relations of molecule abundances with cloud age, and of CO intensity and X_{CO} factor with dust extinction, are insensitive to both the pressure floor and the resolution, although there is more scatter in these relations when we use a lower pressure floor or higher resolution.

A final caveat to note is that these conclusions depend on how the cloud age is defined. We have defined the cloud age from the time when half of the particles currently in the cloud were in a cloud progenitor. However, as we discussed in section 3.2, alternative definitions can result in different ages. Furthermore, our definition requires that we trace individual gas particles back in time. This is trivial in SPH simulations, but is not possible in observations, for which we only have a single snapshot of the cloud at the present day. Observational estimates of GMC ages and lifetimes typically use nearby signatures of star formation, such as young stellar clusters and HII regions (e.g. [Kawamura et al. 2009](#)). We therefore need to be careful when comparing cloud ages from our simulations with observational estimates, as the two definitions may not be equivalent.

ACKNOWLEDGMENTS

We are very grateful to Volker Springel for sharing GADGET3 and his initial conditions code, to Claudio Dalla Vecchia for allowing us to use ANARCHY, and to Benjamin Oppenheimer for code contributions. We thank Simon Glover and Ewine van Dishoeck for useful discussions. We gratefully acknowledge support from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Grant agreement 278594-GasAroundGalaxies. This work used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grant ST/H008519/1, and STFC DiRAC Operations grant ST/K003267/1 and Durham University. DiRAC is part of the National E-Infrastructure. This work also used computer resources provided by the Gauss Centre for Supercomputing/Leibniz Supercomputing Centre under grant:pr83le. We further acknowledge PRACE for awarding us access to resource Supermuc based in Germany at LRZ Garching (proposal number 2013091919). AJR is supported by the Lindheimer Fellowship at Northwestern University.

REFERENCES

- Abdo A. A. et al., 2010, *ApJ*, 710, 133
 Ackermann M. et al., 2010, *A&A*, 538, 71
 Bash F. N., Green E., Peters W. L., III, 1977, *ApJ*, 217, 464
 Bate M. R., Burkert A., 1997, *MNRAS*, 288, 1060
 Bell T. A., Roueff E., Viti S., Williams D. A., 2006, *MNRAS*, 371, 1865
 Bertoldi F., McKee C. F., 1992, *ApJ*, 395, 140
 Black J. H., 1987, *ASSL*, 134, 731
 Bolatto A. D., Leroy A. K., Rosolowsky E., Walter F., Blitz L., 2008, *ApJ*, 686, 948
 Bolatto A. D., Wolfire M., Leroy A. K., 2013, *ARA&A*, 51, 207
 Booth C. M., Theuns T., Okamoto T., 2007, *MNRAS*, 376, 1588
 Cazaux S., Tielens A. G. G. M., 2002, *ApJ*, 575, L29
 Chabrier G., 2003, *PASP*, 115, 763
 Clark P. C., Glover S. C. O., 2015, *MNRAS*, 452, 2057
 Clemens D. P., Sanders D. B., Scoville N. Z., Solomon P. M., 1986, *ApJS*, 60, 297
 Cullen L., Dehnen W., 2010, *MNRAS*, 408, 669
 Dabrowski I., 1984, *Can. J. Phys.*, 62, 1639
 Dale J. E., Bonnell I. A., Clarke C. J., Bate M. R., 2005, *MNRAS*, 358, 291
 Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 426, 140
 Dame T. M., Hartmann D., Thaddeus P., 2001, *ApJ*, 547, 792
 Dobbs C. L., Burkert A., Pringle J. E., 2011, *MNRAS*, 413, 2935
 Dobbs C. L., Pringle J. E., 2013, *MNRAS*, 432, 653
 Dobbs C. L., Pringle J. E., Duarte-Cabral A., 2015, *MNRAS*, 446, 3608
 Duarte-Cabral A., Dobbs C. L., 2016, *arXiv:1603.02470*
 Duffy A. R., Schaye J., Kay S. T., Dalla Vecchia C., 2008, *MNRAS*, 390, 64
 Durier F., Dalla Vecchia C., 2012, *MNRAS*, 419, 465
 Elmegreen B. G., 2000, *ApJ*, 530, 277
 Feldmann R., Gnedin N. Y., Kravtsov A. V., 2012, *ApJ*, 747, 124
 Ferland G. J. et al., 2013, *RMxAA*, 49, 137
 Frerking M. A., Langer W. D., Wilson R. W., 1982, *ApJ*, 262, 590
 Glover S. C. O., Mac Low M.-M., 2007a, *ApJS*, 169, 239
 Glover S. C. O., Mac Low M.-M., 2007b, *ApJ*, 659, 1317
 Glover S. C. O., Mac Low M.-M., 2011, *MNRAS*, 412, 337
 Glover S. C. O., Clark P. C., 2012, *MNRAS*, 426, 377
 Glover S. C. O., Clark P. C., 2016, *MNRAS*, 456, 3596
 Gnedin N. Y., Tassis K., Kravtsov A. V., 2009, *ApJ*, 697, 55
 Gnedin N. Y., Kravtsov A. V., 2010, *ApJ*, 728, 88
 Gray W. J., Scannapieco E., Kasen D., 2015, *ApJ*, 801, 107
 Haardt F., Madau P., 2001, in Neumann D. M., Tran J. T. V., eds, *XXIst Moriond Astrophysics Meeting. Clusters of Galaxies and the High Redshift Universe Observed in X-rays* Editions Frontieres, Paris, 64
 Hernquist L., 1990, *ApJ*, 356, 359
 Heyer M. H., Carpenter J. M., Snell R. L., 2001, *ApJ*, 551, 852
 Heyer M., Krawczyk C., Duval J., Jackson J. M., 2009, *ApJ*, 699, 1092
 Hopkins P. F., Quataert E., Murray N., 2011, *MNRAS*, 417, 950
 Hopkins P. F., Quataert E., Murray N., 2012, *MNRAS*, 421, 3488
 Hopkins P. F., 2013, *MNRAS*, 428, 2840
 Hu C.-Y., Naab T., Walch S., Glover S. C. O., Clark P. C., 2016, *MNRAS*, 458, 3528
 Jackson J. M. et al., 2006, *ApJS*, 163, 145
 Kawamura et al., 2009, *ApJS*, 184, 1
 Kravtsov A. V., 2003, *ApJ*, 590, L1
 Krumholz M. R., McKee C. F., Tumlinson J., 2008, *ApJ*, 689, 865
 Krumholz M. R., Matzner C. D., 2009, *ApJ*, 703, 1352
 Krumholz M. R., McKee C. F., Tumlinson J., 2009, *ApJ*, 693, 216
 Krumholz M. R., Gnedin N. Y., 2011, *ApJ*, 729, 36
 Larson R. B., 1981, *MNRAS*, 194, 809
 Lee C., Leroy A. K., Schnee S., Wong T., Bolatto A. D., Indebetouw R., Rubio M., 2015, *MNRAS*, 450, 2708
 Le Petit F., Nehmé C., Le Bourlot J., Roueff E., 2006, *ApJS*, 164, 506
 Leroy A. K. et al., 2011, *ApJ*, 737, 12
 Lombardi M., Alves J., Lada C. J., 2006, *A&A*, 454, 781
 Martin P. G., Whittet D. C. B., 1990, *ApJ*, 357, 113
 Mathis J. S., Rimpl W., Nordsieck K. H., 1977, *ApJ*, 217, 425
 McKee C. F., Ostriker E. C., 2007, *ARA&A*, 45, 565

McKee C. F., Krumholz M. R., 2010, *ApJ*, 709, 308
Micic M., Glover S. C. O., Federrath C., Klessen R. S., 2012, *MNRAS*, 421, 2531
Miura R. E. et al., 2012, *ApJ*, 761, 37
Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
Murray N., 2011, *ApJ*, 729, 133
Narayanan D., Krumholz M., Ostriker E. C., Hernquist L., 2011, *MNRAS*, 418, 664
Narayanan D., Krumholz M. R., Ostriker E. C., Hernquist L., 2012, *MNRAS*, 421, 3127
Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
Nelson R. P., Langer W. D., 1997, *ApJ*, 482, 796
Oka T., Hasegawa T., Sato F., Tsuboi M., Miyazaki A., Sugimoto M., 2001, *ApJ*, 562, 348
Oppenheimer B. D., Schaye J., 2013b, *MNRAS*, 434, 1063
Pan. H.-A., Fujimoto Y., Tasker E. J., Rosolowsky E., Colombo D., Benincasa S. M., Wadsley J., 2016, *MNRAS*, 458, 2443
Pineda J. E., Caselli P., Goodman A. A., 2008, *ApJ*, 679, 481
Planck Collaboration XIX, 2011, *A&A*, 536, 19
Portinari L., Chiosi C., Bressan A., 1998, *A&A*, 334, 505
Price D. J., 2008, *J. Comput. Phys.*, 227, 10040
Rémy-Ruyer A. et al., 2014, *A&A*, 563, A31
Richings A. J., Schaye J., Oppenheimer B. D., 2014a, *MNRAS*, 440, 3349
Richings A. J., Schaye J., Oppenheimer B. D., 2014b, *MNRAS*, 442, 2780
Richings A. J., Schaye J., 2016, *MNRAS*, 458, 270
Robertson B. E., Kravtsov A. V., 2008, *ApJ*, 680, 1083
Roman-Duval J., Jackson J. M., Heyer M., Rathborne J., Simon R., 2010, *ApJ*, 723, 492
Rosdahl J., Schaye J., Teyssier R., Agertz O., 2015, *MNRAS*, 451, 4553
Rosolowsky E., 2007, *ApJ*, 654, 240
Sanders D. B., Clemens D. P., Scoville N. Z., Solomon P. M., 1986, *ApJS*, 60, 1
Sawala T. et al., 2015, *MNRAS*, 448, 2941
Schaye J., 2001, *ApJ*, 562, L95
Schaye J., Dalla Vecchia C., 2008, *MNRAS*, 383, 1210
Schaye J. et al., 2015, *MNRAS*, 446, 521
Schöier F. L., van der Tak F. F. S., van Dishoeck E. F., Black J. H., 2005, *A&A*, 432, 369
Scoville N. Z., Solomon P. M., Sanders D. B., 1979, in Burton W. D., ed., *Proc. IAU Symp. 84, The Large-Scale Characteristics of the Galaxy*, Reidel, Dordrecht, p. 277
Scoville N. Z., Yun M. S., Clemens D. P., Sanders D. B., Waller W. H., 1987, *ApJS*, 63, 821
Shetty R., Glover S. C., Dullemond C. P., Klessen R. S., 2011a, *MNRAS*, 412, 1686
Shetty R., Glover S. C., Dullemond C. P., Ostriker E. C., Harris A. I., Klessen R. S., 2011b, *MNRAS*, 415, 3253
Smith R. J., Glover S. C. O., Clark P. C., Klessen R. S., Springel V., 2014, *MNRAS*, 441, 1628
Sobolev V. V., 1957, *SvA*, 1, 678
Solomon P. M., Rivolo A. R., Barrett J., Yahil A., 1987, *ApJ*, 319, 730
Springel V., 2005, *MNRAS*, 364, 1105
Springel V., Di Matteo T., Hernquist L., 2005, *MNRAS*, 361, 776
Springel V., 2010, *MNRAS*, 401, 791
Sternberg A., Le Petit F., Roueff E., Le Bourlot J., 2014, *ApJ*, 190, 10
Strong A. W., Mattox J. R., 1996, *A&A*, 308, L21
Tielens A. G. G. M., Hollenbach D., 1985, *ApJ*, 291, 722
Truelove J. K., Klein R. I., McKee C. F., Holliman J. H. II, Howell L. H., Greenough J. A., 1997, *ApJ*, 489, L179
van Dishoeck E. F., Black J. H., 1988, *ApJ*, 334, 771
van Dishoeck E. F., Jonkheid B., van Hemert M. C., 2006, *Faraday Discuss.*, 133, 231
Visser R., van Dishoeck E. F., Black J. H., 2009, *A&A*, 503, 323




Figure A1. Cloud mass functions for the ref model with a mass of $187.5 M_{\odot}$ (*hiRes*; red dashed curve), $750 M_{\odot}$ (*ref*; black solid curve) and $3000 M_{\odot}$ (*lowRes*; blue dot-dashed curve) per SPH particle. The vertical dotted lines indicate the cloud mass containing 50 particles for each resolution. We include snapshots from 100 to 400 Myr at 100 Myr intervals. The mass functions of the ref and hiRes runs agree at masses $\gtrsim 10^4 M_{\odot}$.

Walch S. K., Whitworth A. P., Bisbas T., Wünsch R., Hubber D., 2012, *MNRAS*, 427, 625
Walch S., Naab T., 2015, *MNRAS*, 451, 2757
Wendland H., 1995, *Adv. Comput. Math.*, 4, 389
Wiersma R. P. C., Schaye J., Smith B. D., 2009, *MNRAS*, 393, 99
Williams J. P., Bergin E. A., Caselli P., Myers P. C., Plume, R., 1998, *ApJ*, 503, 689
Yang B., Stancil P. C., Balakrishnan N., Forrey R. C., 2010, *ApJ*, 718, 1062

APPENDIX A: RESOLUTION TESTS

In this paper we have focussed on the properties of GMCs in our simulations of isolated galaxies. These GMCs are among the smallest structures that we can resolve in these simulations, and thus are close to our resolution limit. We therefore need to test how sensitive our results are to the resolution of our simulations.

We repeated our ref model ($Z = 0.1 Z_{\odot}$ and the ISRF of Black 1987) twice, with a factor four higher and lower mass resolution (runs hiRes and lowRes, respectively). We also decreased/increased the gravitational softening by a factor $4^{1/3}$, with $\epsilon_{\text{soft}} = 2.0$ pc and 5.0 pc in the hiRes and lowRes runs, respectively. For these resolution tests, we set the pressure floor such that $N_{J,m} = 4$, as for our fiducial resolution (see section 2.1). See Appendix B for the effects of varying the pressure floor at fixed resolution. We ran the hiRes and lowRes simulations for 500 Myr, then we identified clumps of gas above a density threshold of $n_{H,\text{min}} = 10 \text{ cm}^{-3}$ using a FoF algorithm as before (see section 3.1), except that we used a FoF linking length of 6.3 pc and 15.9 pc in the hiRes and lowRes runs, respectively. These linking lengths correspond to the mean interparticle spacing at the density threshold for the given resolution.

Fig A1 shows the cloud mass functions for the three resolutions. The vertical dotted lines indicate the cloud mass




Figure A2. Cloud age, defined using the particles originally in the cloud in the current snapshot (see section 3.2), versus mass for clouds in the ref model with a mass of $187.5 M_\odot$, $750 M_\odot$ and $3000 M_\odot$ per SPH particle (left, centre and right panels, respectively). We include all clouds with at least 50 particles in snapshots at 100 Myr intervals from 100 to 400 Myr. The horizontal stripes of points at low ages are due to the finite interval (1 Myr) between snapshots. The cloud ages do not depend strongly on the resolution of the simulation.




Figure A3. Mass-size relation for the ref model with three different particle masses: $187.5 M_\odot$, $750 M_\odot$ and $3000 M_\odot$ (left, centre and right panels, respectively). The colour scale indicates the cloud age. We include all clouds with at least 50 particles in snapshots at 100 Myr intervals from 100 to 400 Myr. We also show the observed relations from Solomon et al. (1987) (S87; solid curves) and Roman-Duval et al. (2010) (RD10; dashed curves). The mass-size relation is insensitive to resolution.

containing 50 particles for each resolution. The mass function is steeper in the lowRes run, but the ref and hiRes runs agree at masses $\gtrsim 10^4 M_\odot$.

Fig. A2 shows the cloud age as a function of cloud mass. The cloud age here uses our fiducial age definition, i.e. the time when half of the particles originally in the cloud in the current snapshot were in a progenitor of the cloud (see section 3.2). We include all clouds in snapshots at intervals of 100 Myr from 100 to 400 Myr. The distribution of cloud ages is not strongly affected by resolution.

Fig. A3 shows the cloud mass-size relation for the three resolutions, from hiRes (left panel) to lowRes (right panel). The colour scale indicates the cloud age, and the solid and dashed lines show the observed relations of Solomon et al. (1987) and Roman-Duval et al. (2010), respectively. All three simulations follow the same relation, with the same slope and normalisation. The hiRes simulation extends this relation to lower cloud masses and sizes, as it can resolve smaller clouds (and we include only clouds with at least 50 particles).

Fig. A4 shows the virial parameter, α , versus cloud




Figure A4. Cloud virial parameter, α , versus mass, for the ref model with three different particle masses: $187.5 M_\odot$, $750 M_\odot$ and $3000 M_\odot$ (left, centre and right panels, respectively). The colour scale indicates the cloud age, and the horizontal dotted lines show a value of $\alpha = 1$, which corresponds to virial equilibrium. As we increase the resolution, the lower envelope of α (which is influenced by the pressure floor) decreases. However, clouds with $M_{\text{cloud}} \gtrsim 10^5 M_\odot$ still have $\alpha > 1$, even in the hiRes run. These high-mass clouds are therefore genuinely unvirialised, and this is not an artifact of the pressure floor.

mass for the three resolutions (different panels). The horizontal dotted line indicates $\alpha = 1$, which corresponds to virial equilibrium. In the ref and hiRes runs we clearly see a lower envelope of $\alpha \propto M^{-2/3}$. As discussed in section 4, this envelope is caused by the pressure floor, which prevents low-mass clouds from becoming gravitationally bound. As we increase the resolution, the envelope decreases, because the Jeans mass of the pressure floor (which we set to four times the kernel mass for these resolution tests) decreases. However, this envelope does not extend below $\alpha = 1$. For example, in the hiRes run, the lower envelope caused by the pressure floor reaches $\alpha = 1$ at $M_{\text{cloud}} \approx 10^5 M_\odot$. Above this mass, the minimum α in our simulated clouds is just above unity, independent of cloud mass, and no longer follows the scaling that would arise from the pressure floor. This suggests that, in our hiRes simulation, clouds with $M_{\text{cloud}} \gtrsim 10^5 M_\odot$ are genuinely unvirialised, with $\alpha > 1$, and this is not an artifact of the pressure floor.

Fig. A5 shows the H_2 fraction, f_{H_2} (top row), and the ratio of f_{H_2} to the equilibrium H_2 fraction, $f_{\text{H}_2}^{\text{eqm}}$ (bottom row), plotted against cloud age, for the three resolutions, from hiRes (left column) to lowRes (right column). The solid black curves show the median relations in bins of age, the horizontal dotted lines indicate a value of unity, and the colour scale indicates the cloud mass. All resolutions follow similar median relations, both for f_{H_2} versus age and $f_{\text{H}_2}/f_{\text{H}_2}^{\text{eqm}}$ versus age. However, there is more scatter at higher resolution. In the hiRes run (left-hand column), the scatter is particularly large for low-mass ($\lesssim 3 \times 10^4 M_\odot$) clouds (dark blue/purple circles), which are not present at lower resolutions because we include only clouds with at least 50 particles.

The CO fractions (not shown) also follow similar median relations with cloud age at different resolutions, albeit with more scatter at higher resolution due to the low-mass clouds.

Fig. A6 shows the mean CO intensity, I_{CO} (top row), and the mean X_{CO} factor (bottom row) versus the mean dust extinction, A_v , for the three resolutions. All three res-




Figure A5. H_2 fraction, f_{H_2} (top row), and the ratio of f_{H_2} to the H_2 fraction in equilibrium, $f_{\text{H}_2}^{\text{eqm}}$, plotted against cloud age, for clouds in the ref model with three different particle masses: $187.5 M_\odot$, $750 M_\odot$ and $3000 M_\odot$ (left, centre and right columns, respectively). The colour scale indicates the cloud mass, the solid curves indicate the median relation in bins of age, and the horizontal dotted lines indicate a value of unity. We include all clouds with at least 50 particles in snapshots at 100 Myr intervals from 100 to 400 Myr. The median relations for the different resolutions agree, although there is more scatter in the hiRes run, which is primarily due to low-mass clouds ($\lesssim 3 \times 10^4 M_\odot$) that are not present at lower resolutions.




Figure A6. Mean CO intensity, I_{CO} (top row), and X_{CO} factor (bottom row), versus mean dust extinction, A_v . The colour scale indicates the cloud age. In the top row, we also show the $I_{\text{CO}} - A_v$ relations observed in the Pipe nebula (Lombardi et al. 2006; solid curves) and the Perseus cloud (Pineda et al. 2008; dot-dashed curves). The horizontal dotted lines in the top row show a value of $I_{\text{CO}} = 0.25 \text{ K km s}^{-1}$, which corresponds to the 3σ intensity threshold for the Small Magellanic Cloud in the observations of Leroy et al. (2011). In the bottom row, the horizontal dotted lines show the typical value measured in the Milky Way, $X_{\text{CO}} = 2 \times 10^{20} \text{ cm}^{-2} (\text{K km s}^{-1})^{-1}$ (e.g. Bolatto et al. 2013). The different resolutions show similar trends of I_{CO} and X_{CO} with A_v .

olutions follow similar trends of I_{CO} and X_{CO} with A_v , and the threshold dust extinction, A_v^{thresh} , below which I_{CO} is strongly suppressed, is insensitive to the resolution.

APPENDIX B: EFFECTS OF THE PRESSURE FLOOR

In our simulations we impose a density-dependent pressure floor (equation 2.1) to ensure that the Jeans mass is always at least a factor $N_{\text{J},\text{m}} = 4$ times the kernel mass. However, we saw in Fig. 10 that this pressure floor prevents the lowest-mass clouds ($\lesssim 3 \times 10^5 M_\odot$) from becoming strongly gravitationally bound ($\alpha < 1$). While the pressure floor could represent sources of pressure that are not explicitly included in our models, such as unresolved turbulence, its functional form was motivated by numerical reasons, so we may overestimate the true pressure.

Conversely, if we remove the pressure floor entirely, so that we no longer resolve the Jeans mass in cold, dense gas, we may experience artificial fragmentation and collapse in gas that should be Jeans-stable. Bate & Burkert (1997) considered the resolution requirements in SPH simulations that include self-gravity, applied to the collapse and fragmentation of molecular clouds. In simulations where the gravitational softening length, ϵ_{soft} , is smaller than the SPH smoothing length, h_{sml} , which is the case for nearly all gas particles in our simulations (we use $\epsilon_{\text{soft}} = 3.1 \text{ pc}$ for the ref model), they found that, if the Jeans mass is not resolved by at least two times the kernel mass, gas can artificially undergo collapse when it should be Jeans-stable. The reason is that the thermal pressure is smoothed on scales of h_{sml} , while the gravitational force is smoothed on scales of ϵ_{soft} . Therefore, once a gas cloud becomes unresolved, the pressure force will be smoothed out before the gravitational force. Hence it will artificially lose pressure support against gravity and will undergo gravitational collapse.

A number of approaches have been proposed in the literature to alleviate these problems of artificial fragmentation that may occur when the Jeans scale is unresolved. Booth et al. (2007) developed a prescription for star formation and feedback in disc galaxies in which unresolved molecular gas is modelled by ‘sticky particles’ that coagulate when they collide. Narayanan et al. (2011) also use a subgrid prescription to model the unresolved cold, molecular gas, which involves hybrid SPH particles that include both the warm and the cold ISM phases. Schaye & Dalla Vecchia (2008) imposed a polytropic equation of state, $P \propto \rho^{\gamma_{\text{eff}}}$, in their SPH simulations of disc galaxies. They used $\gamma_{\text{eff}} = 4/3$ so that $m_{\text{gas}}/M_{\text{J}}$ and $h_{\text{sml}}/L_{\text{J}}$ are constant, where m_{gas} is the SPH particle mass, h_{sml} is the SPH smoothing length, and M_{J} and L_{J} are the Jeans mass and length, respectively. They set the normalisation of the equation of state such that $N_{\text{J},\text{m}} = 6$ in our nomenclature. Robertson & Kravtsov (2008) imposed a pressure floor in their SPH simulations of disc galaxies by setting a minimum internal energy (i.e. a temperature floor). The functional form of their temperature floor was equivalent to a polytropic equation of state with $P \propto \rho^{4/3}$. They explored a range of temperature floors, corresponding to $N_{\text{J},\text{m}} = 2 - 200$ (see their appendix A), and they used a fiducial floor with $N_{\text{J},\text{m}} = 30$. Hopkins et al. (2011) used the pressure floor of Robert-




Figure B1. Cloud mass functions for the ref model, using a pressure floor corresponding to a constant Jeans mass with $N_{J,m} = 4$ (black solid curve), $N_{J,m} = 1$ (red dashed curve) and $N_{J,m} = 0.25$ (blue dot-dashed curve), where $N_{J,m}$ is the ratio of the minimum Jeans mass to the mass within the SPH kernel. The cloud mass function is insensitive to the pressure floor.

son & Kravtsov (2008) in their SPH simulations of disk galaxies, with $N_{J,m} = 10$ (although they also considered $N_{J,m} = 4 - 15$, see their appendix A).

However, not all simulation studies have imposed a Jeans limiter. For example, Clark & Glover (2015) simply ended their SPH simulations of molecular clouds once the Jeans scale became unresolved. However, they used a very high resolution ($0.005 M_\odot$ per particle and 50 SPH neighbours), so they could follow the gravitational collapse up to a density of $\sim 10^6 \text{ cm}^{-3}$. Walch (2015) also did not include a Jeans limiter in their SPH simulations of supernova feedback in molecular clouds. The densities in some of their simulations extended up to $\sim 10^6 \text{ cm}^{-3}$ (see their fig. 3), but they used a lower resolution than Clark & Glover (2015), with $0.1 M_\odot$ per SPH particle, so the Jeans scale will be unresolved in some of their simulations. Hu et al. (2016) also did not include a Jeans limiter in their simulations of dwarf galaxies, although they did check that the resolution of their simulations ($4 M_\odot$ per SPH particle, with 100 neighbours) is high enough that the Jeans scale is resolved in the majority of the gas. There are also other studies that do not explicitly include a Jeans limiter, for example Gnedin & Kravtsov (2011); Glover & Mac Low (2011); Shetty et al. (2011a,b); Dobbs & Pringle (2013).

To explore how the pressure floor affects our results, we repeated the ref model twice, with the pressure floor reduced by factors of 4 and 16 in terms of the Jeans mass (i.e. with $N_{J,m} = 1$ and 0.25 , respectively), using our fiducial resolution of $750 M_\odot$ per SPH particle. These simulations were only run for 100 Myr, and we therefore compared clouds at 100 Myr in these runs with the ref model, whereas previously we have been combining clouds from nine snapshots (or four snapshots for our resolution tests in Appendix A) at 100 Myr intervals.

Fig. B1 shows the cloud mass functions, which are similar for the three pressure floors. In Fig. B2 we show the cloud age as a function of mass, where the cloud age is de-




Figure B2. Cloud age, defined using the particles originally in the cloud in the current snapshot (see section 3.2), versus mass for clouds in the snapshot at 100 Myr in the ref model with three different pressure floors: $N_{J,m} = 4$, 1 and 0.25 (left, centre and right panels, respectively). In the run with the lowest pressure floor (right panel), low-mass, poorly resolved clouds can survive for longer (up to 50 Myr) than in the run with our fiducial pressure floor (left panel).




Figure B3. Mass-size relation for clouds in the snapshot at 100 Myr, from the ref model with three different pressure floors: $N_{J,m} = 4$, 1 and 0.25 (left, centre and right panels, respectively). The colour scale indicates the cloud age. We also show the observed relations from Solomon et al. (1987) (S87; solid curves) and Roman-Duval et al. (2010) (RD10; dashed curves). In the run with $N_{J,m} = 0.25$, low-mass, poorly resolved clouds are more compact, but this run no longer reproduces the observed slope of this relation.

fined by the time when half of the particles that were originally in the cloud in the current snapshot (100 Myr) were in a progenitor of that cloud, i.e. our fiducial age definition (see section 3.2). The left, centre and right panels show the results for a pressure floor with $N_{J,m} = 4$, 1 and 0.25, respectively. In the run with the lowest pressure floor (right panel), we find clouds extending up to higher ages (up to 50 Myr). However, the longest-lived clouds ($\gtrsim 30$ Myr) in this simulation are only found in the lowest-mass, most poorly resolved clouds ($M_{\text{cloud}} \lesssim 10^5 M_\odot$), whereas we see no trends between cloud age and mass with higher pressure floors.

Fig. B3 shows the mass-size relation for clouds with the different pressure floors. The colour scale indicates the cloud age, and the black solid and dashed curves show the observed relations of Solomon et al. (1987) and Roman-Duval et al. (2010), respectively. The run with the lowest pressure floor (right panel) shows more scatter in this relation, with clouds of the same mass generally being more compact than in the $N_{J,m} = 4$ run. Additionally, the most compact clouds, i.e. those that lie to the left of the relation, tend to be older (\gtrsim




Figure B4. Cloud virial parameter, α , versus mass, for the ref model with three different pressure floors: $N_{J,m} = 4, 1$ and 0.25 (left, centre and right panels, respectively). The colour scale indicates the cloud age, and the horizontal dotted lines show a value of $\alpha = 1$, which corresponds to virial equilibrium. The lower envelope of α decreases as we lower the pressure floor, with some low-mass clouds in the $N_{J,m} = 0.25$ run becoming strongly gravitationally bound. However, it is likely that the most compact clouds in the $N_{J,m} = 0.25$ run are affected by spurious fragmentation and collapse that may arise when we do not fully resolve the Jeans scale (see text).

30 Myr). This is understandable, as more compact objects will be more strongly gravitationally bound, and thus can survive for longer. This explains the longer cloud ages that we saw in Fig. B2 in the $N_{J,m} = 0.25$ run.

However, it is only the lowest-mass, most poorly resolved clouds that are more compact in the $N_{J,m} = 0.25$ run than in the $N_{J,m} = 4$ run. The highest-mass clouds ($\gtrsim 3 \times 10^5 M_\odot$) appear unaffected by the pressure floor. This also means that the mass-size relation in the $N_{J,m} = 0.25$ run is flatter than is observed. While the clouds in the $N_{J,m} = 4$ run lie further from the observed relations, they do recover the same slope, and the difference in normalisation can be explained by the differences in cloud definition (as seen in Figs. 6 and 7). Furthermore, in the hiRes run in Appendix A, the mass-size relation did not change when we increased the resolution, and we did not see low-mass clouds become more compact with longer ages. This suggests that the more compact clouds that we find when we lower the pressure floor may be an artifact of spurious fragmentation and collapse that may arise when we do not fully resolve the Jeans scale.

Fig. B4 shows the virial parameter, α , as a function of cloud mass, for the three pressure floors (in different panels). The horizontal dotted line indicates a value of $\alpha = 1$, which corresponds to virial equilibrium. In the left panel, for our fiducial pressure floor, $N_{J,m} = 4$, we see a lower envelope in α that declines with cloud mass as $\alpha \propto M_{\text{cloud}}^{-2/3}$. This scaling is due to the pressure floor, as discussed in section 4, and means that, in the $N_{J,m} = 4$ run, the pressure floor prevents the lowest-mass clouds from becoming strongly gravitationally bound ($\alpha < 1$).

As we lower the pressure floor, the low-mass clouds extend to lower values of α , and can become strongly gravitationally bound ($\alpha < 1$) in the $N_{J,m} = 0.25$ run (although there are still clouds with high values of α as well). This is consistent with the more compact low-mass clouds that we saw in this run in Fig. B3. However, as noted above, the more compact low-mass clouds that we find when we lower




Figure B5. H_2 fraction, f_{H_2} (top row), and the ratio of f_{H_2} to the H_2 fraction in equilibrium, $f_{H_2}^{\text{eqm}}$, plotted against cloud age, for clouds in the snapshot at 100 Myr from the ref model with three different pressure floors: $N_{J,m} = 4, 1$ and 0.25 (left, centre and right columns, respectively). The colour scale indicates the cloud mass, the solid curves indicate the median relation in bins of age, and the horizontal dotted lines indicate a value of unity. As we lower the pressure floor, clouds extend to higher ages, but continue to follow similar median relations of f_{H_2} and $f_{H_2}/f_{H_2}^{\text{eqm}}$ with age, albeit with more scatter.

the pressure floor are likely to be an artifact of spurious fragmentation and collapse. Furthermore, there were no clouds with $\alpha < 1$ in the hiRes run in Appendix A.

We now need to consider whether varying the pressure floor affects the evolution of molecular abundances, the CO emission or the X_{CO} factors of individual clouds. Fig. B5 shows the H_2 fraction, f_{H_2} (top row), and the ratio of f_{H_2} to the equilibrium H_2 fraction, $f_{H_2}^{\text{eqm}}$ (bottom row), plotted against cloud age, for the three pressure floors (in different columns). As we lower the pressure floor (left to right), clouds extend to higher ages, as we saw in Fig. B2. However, the median values of f_{H_2} and $f_{H_2}/f_{H_2}^{\text{eqm}}$ at fixed cloud age are similar, regardless of the pressure floor. Thus, by lowering the pressure floor, we simply extend the same median relations of f_{H_2} and $f_{H_2}/f_{H_2}^{\text{eqm}}$ versus cloud age to higher ages, although the scatter also increases. In particular, in all three runs it takes $\approx 10 - 15$ Myr for the median f_{H_2} to reach within a factor two of its equilibrium value.

When we lower the pressure floor, the CO fractions (not shown) also follow similar median relations with cloud age as for our fiducial pressure floor, albeit with more scatter.

Fig. B6 shows the mean CO intensity, I_{CO} (top row), and the mean X_{CO} factor (bottom row) of each cloud for the three pressure floors as a function of mean dust extinction, A_v . As we lower the pressure floor, clouds extend to higher A_v , as they become more compact. However, they continue to follow a similar trend of I_{CO} with A_v , although there is more scatter in this relation with a lower pressure floor. In particular, the threshold A_v below which I_{CO} cuts off does not appear to be strongly affected by the pressure floor, i.e. the $I_{\text{CO}} - A_v$ relation does not move horizontally in this plot. However, the increased scatter as we reduce the pressure floor means that we find some low- A_v clouds ($A_v \approx 0.06$)




Figure B6. Mean CO intensity, I_{CO} (top row), and X_{CO} factor (bottom row), plotted against mean dust extinction, A_v . The colour scale indicates the cloud age. In the top row, we also show the $I_{\text{CO}} - A_v$ relations observed in the Pipe nebula (Lombardi et al. 2006; solid curves) and the Perseus cloud (Pineda et al. 2008; dot-dashed curves). The horizontal dotted lines in the top row indicate a value of $I_{\text{CO}} = 0.25 \text{ K km s}^{-1}$, which corresponds to the 3σ intensity threshold for the Small Magellanic Cloud in the observations of Leroy et al. (2011). In the bottom row, the horizontal dotted lines indicate the typical value measured in the Milky Way, $X_{\text{CO}} = 2 \times 10^{20} \text{ cm}^{-2} (\text{K km s}^{-1})^{-1}$ (e.g. Bolatto et al. 2013). As we lower the pressure floor, clouds extend to higher A_v , as they become more compact. However, for the values of A_v where the different runs overlap, the $I_{\text{CO}} - A_v$ relations are consistent, albeit with more scatter, particularly in I_{CO} , and with a shallower cut-off in I_{CO} towards low A_v .

with higher CO intensities ($I_{\text{CO}} \approx 0.3 \text{ K km s}^{-1}$), so the cut-off in I_{CO} at low A_v is less steep than in the $N_{\text{J},m} = 4$ run.

The X_{CO} factors all show similar scatter of four orders of magnitude, regardless of pressure floor. In the $N_{\text{J},m} = 0.25$ run, the high- A_v clouds ($\gtrsim 0.3$) show less scatter in X_{CO} (one order of magnitude) and suggest a trend of decreasing X_{CO} with increasing A_v . This trend cannot be verified in the $N_{\text{J},m} = 4$ run, because it does not include these high- A_v clouds. However, for the range of A_v where these three runs overlap, they show consistent $X_{\text{CO}} - A_v$ relations.

To conclude, lowering the pressure floor results in low mass clouds ($\lesssim 3 \times 10^5 M_\odot$) becoming more compact, more strongly gravitationally bound, and able to survive for longer. However, this is inconsistent with what we found when we increased the resolution in Appendix A, so it is likely that this is an artifact of spurious fragmentation and collapse that may arise when we do not fully resolve the Jeans scale. In particular, the run with the lowest pressure floor ($N_{\text{J},m} = 0.25$) no longer reproduces the observed slope of the cloud mass-size relation (Solomon et al. 1987; Roman-Duval et al. 2010), as it becomes flatter when the pressure floor is lowered, as the low-mass clouds become too compact. This flattening of the mass-size relation may be a characteristic signature that we can use in the future to determine whether a given simulation suffers from spurious fragmentation due to an unresolved Jeans scale.

Furthermore, we find similar median trends of f_{H_2} and

$f_{\text{H}_2}/f_{\text{H}_2}^{\text{eqm}}$ versus cloud age, and I_{CO} and X_{CO} versus A_v , regardless of pressure floor. Lowering the pressure floor simply extends these relations to higher ages and A_v , albeit with more scatter.